

Systematic Generation of Algorithms for Iterative Methods

Master's Thesis

Henrik Barthels, B.Sc.

Supervised by
Prof. Paolo Bientinesi, Ph.D.
Prof. Georg May, Ph.D.

Revised Edition

This is a revised edition of the author's thesis. It differs from the original version in corrections of typographical errors and clarifications of some passages.

Abstract

The FLAME methodology makes it possible to derive provably correct algorithms from a formal description of a linear algebra problem. So far, the methodology has been successfully used to automate the derivation of direct algorithms such as the Cholesky decomposition and the solution of Sylvester equations. In this thesis, we present an extension of the FLAME methodology to tackle iterative methods such as Conjugate Gradient. As a starting point, we use a formal description of the iterative method in matrix form. The result is a family of provably correct pseudocode algorithms. We argue that all the intermediate steps are sufficiently systematic to be fully automated.

Acknowledgments

I would like to thank Paolo Bientinesi and Diego Fabregat Traver for their support, guidance, and taking their time for all those lengthy discussions. I want to thank Georg May for agreeing to be the second supervisor of this thesis. Last but not least, I thank Friederike and Robert for proofreading.

Contents

1. Introduction	1
1.1. Background: FLAME	1
1.2. Challenges	2
1.3. Contributions	3
1.4. Outline of the Thesis	4
2. Derivation of Algorithms for Direct Methods	5
2.1. Triangular Linear System	5
2.2. Symmetric Positive Definite Linear System	11
2.3. General Linear System	15
2.4. Equivalence of Loop Invariants	16
3. Derivation of Algorithms for Iterative Methods: Foundations	23
3.1. Matrix Representation	23
3.2. Systematic Derivation of Matrix Properties	25
3.2.1. Preliminaries	25
3.2.2. Representing Knowledge about Linear Algebra	26
3.2.3. Derivation of Properties	27
3.2.4. Design Considerations	29
3.2.5. Example: Nonsymmetric CG	31
3.3. Initial Partitionings	35
4. Derivation of Algorithms for Iterative Methods	43
4.1. Derivation of Algorithms	43
4.1.1. PME Generation	43
4.1.2. Loop Invariant Identification	46
4.1.3. Algorithm Construction	51
4.1.4. Refinement	53
4.2. Example: Nonsymmetric CG	55
4.3. Example: BiCG	60
4.3.1. Remark on the Equivalence of Loop Invariants and Algorithms	69
4.4. Scope and Limitations	69
5. Conclusion	73
A. Matrix Properties	75

B. Matrix Representations of Iterative Methods	77
B.1. Krylov Subspace Methods	77
B.1.1. Krylov Sequence	77
B.1.2. Steepest Descent	77
B.1.3. Conjugate Gradient (symmetric)	78
B.1.4. Conjugate Gradient (nonsymmetric)	78
B.1.5. BiCG	79
B.2. Stationary Iterative Methods	79
B.2.1. Gauss-Seidel Method	79
B.2.2. Jacobi Method	80
B.2.3. Successive Overrelaxation Method	80
B.2.4. Richardson Iteration	80
Bibliography	81

Glossary

A, B, C, \dots	Matrices.
a, b, c, \dots	Row or column vectors.
$\alpha, \beta, \gamma, \dots$	Scalars.
BiCG	Biconjugate Gradient method.
BLAS	Basic Linear Algebra Subprograms. A library for linear algebra operations.
CG	Conjugate Gradient method.
e_0	The unit vector that is one in the first position.
e_r	The unit vector that is one in the last position.
FLAME	Formal Linear Algebra Methods Environment.
G	Loop guard.
I	Identity matrix.
P_{after}	After predicate.
P_{before}	Before predicate.
P_{inv}	Loop invariant.
PME	Partitioned Matrix Expression.
P_{post}	Postcondition.
P_{pre}	Precondition.
SPD	Symmetric Positive Definite.
$\ \dots\ $	An arbitrary vector norm.

1. Introduction

The goal of this thesis is to simplify the development of algorithms for iterative methods. We present a methodology for the systematic derivation of such algorithms and lay the foundations for a system that automates the generation of algorithms and code for iterative solvers for linear systems.

Over the last few decades, iterative methods have become an indispensable tool for solving sparse linear systems. Such systems commonly occur in science and engineering, for instance when discretized partial differential equations have to be solved. While direct methods are a reliable tool to solve linear systems, their ability to use the sparsity of a matrix to their advantage is limited to specific sparsity patterns. Often enough, those sparse systems are so large that using direct methods is impractical.

Since the introduction of the Conjugate Gradient (CG) method in 1952 [9], much progress has been made in the field (for an overview, consider [1, 11, 13]). However, the way from an expert's idea for an algorithm to a working implementation is still a long one. At first, a formal description of the algorithm has to be derived on paper. Then, it has to be shown that the algorithm is correct, something that ideally follows from the derivation. Furthermore, to assess how useful the algorithm is in practice, it is desirable to prove that it is numerically stable. Finally, the algorithm has to be translated into code. Usually, that is done multiple times, for different languages, or potentially using different libraries. Every one of those steps takes time and is a potential source of errors.

To speed up this process and eliminate those sources of error, one may try to automate some or all of the steps described above. Automating the ingenuity of the iterative methods expert certainly lies in the distant future. In contrast, automatically translating a sufficiently formal description of an algorithm into code is a lot more feasible.

This thesis covers the systematical derivation of provably correct (pseudocode) algorithms, based on an abstract, formal description of an iterative method.

1.1. Background: FLAME

The Formal Linear Algebra Methods Environment (FLAME) is a project with the goal to automate the derivation of linear algebra algorithms, as well as their implementation. In [2], it was shown that it is possible to systematically derive algorithms for dense linear algebra in a number of well defined steps.

1. Introduction

The starting point of this derivation is a formal description of the input, consisting of a precondition (P_{pre}) and a postcondition (P_{post}). As an example, the description of a linear system where A is lower triangular is shown below:

$$x := \Phi(A, b) \equiv \begin{cases} P_{\text{pre}} : \{\text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\ \quad \text{LowerTriangular}[A] \wedge \\ \quad \text{Input}[b] \wedge \text{Vector}[b] \wedge \\ \quad \text{Output}[x] \wedge \text{Vector}[x]\} \\ P_{\text{post}} : \{Ax = b\} \end{cases}$$

The function $x := \Phi(A, b)$ is used to abstract from the details of this representation.

At the core of the derivation lies the concept of a Partitioned Matrix Expression (PME). A PME describes all parts of the output operands of a linear algebra operation in terms of parts of the input operands. For the lower triangular system, the equation $Ax = b$ is partitioned as

$$\left(\begin{array}{c|c} A_{\text{TL}} & 0 \\ \hline A_{\text{BL}} & A_{\text{BR}} \end{array} \right) \begin{pmatrix} x_{\text{T}} \\ x_{\text{B}} \end{pmatrix} = \begin{pmatrix} b_{\text{T}} \\ b_{\text{B}} \end{pmatrix},$$

where A_{TL} and A_{BR} are square. Then, the PME is

$$\begin{pmatrix} x_{\text{T}} := \Phi(A_{\text{TL}}, b_{\text{T}}) \\ x_{\text{B}} := \Phi(A_{\text{BR}}, b_{\text{B}} - A_{\text{BL}}x_{\text{T}}) \end{pmatrix}.$$

The approach described in [2] has a number of advantages: (1) The resulting algorithms are built around a proof of correctness, so they are correct by construction. (2) It naturally leads to multiple variants of algorithms for the same operation. While they are all correct in exact arithmetic, they potentially behave very different in presence of round-off errors. (3) Furthermore, it was shown that this derivation can be combined with a systematic stability analysis.

The author of [2] presented evidence that this derivation, based on a formal description of the operation and the PME, is systematic enough to be automated. A system that automates this process, including the generation of PMEs, was presented in [6].

All of the efforts described above focused on direct methods. Naturally, the question arises if this methodology extends to iterative methods. In [5], Eijkhout et al. introduced a matrix representation of the CG method and showed preliminary evidence that a systematic derivation of algorithms with a FLAME-like methodology is possible. However, the approach presented there heavily relied on guidance by a human expert. This thesis can be seen as a continuation of this work.

1.2. Challenges

The systematic derivation of algorithms for iterative methods introduces new challenges, especially for deriving the PME.

- For direct methods, the sizes of all operands are constant. In case of iterative methods, a variable number of iterations can be performed. In the matrix representation, this is reflected by the fact that some operands have variable sizes.
- The matrix representation of iterative methods introduces new types of operands. Usually, an operand is either known or unknown, so it is input or output, respectively. Now, there are operands that are initially partially known and partially unknown.
- Quite often, equations have to be solved by using properties of certain expressions. Consider the following equation as an example:

$$-Pu + p = r$$

P , p and r are known, and the goal is to find an assignment for u . It can be solved by using the fact that P^TAP is lower triangular and P^TAp is zero (A is known as well). Multiplying P^TA from the left to both sides of the equation results in

$$-P^TAPu = P^TA r,$$

which is a triangular system that can easily be solved.

This introduces two challenges: The properties of P^TAP and P^TAp are not explicitly part of the matrix representation, so they have to be derived from it by algebraic manipulations and deductive reasoning. Then, to enable a system to solve this equation, it must be capable of identifying that the properties of those expressions can be used to do so.

- To ensure that the derived algorithms can be used in practice, it is desirable to derive the exact same algorithms as used today. The difficulty in achieving this goal lies in the fact that some of those algorithms are the result of nontrivial rewritings of easily derivable formulas. Those rewritings are hard to automate.

1.3. Contributions

The following contributions are made in this thesis:

Systematic derivation of algorithms for iterative methods. We present an approach for systematically deriving algorithms for iterative methods similar to the one discussed above for direct methods. To achieve this, we present solutions to the following two problems:

Systematic derivation of matrix properties. To be able to derive algorithms, properties of matrices or matrix expressions are needed that are not explicitly part of the formal description of the operation. We present a method

1. Introduction

for deriving those properties from the description, using a set of inference rules that encode linear algebra knowledge.

Systematic derivation of PME for iterative methods. The derivation of PME for iterative methods is more complex compared to direct methods. As part of this derivation, equations have to be solved using the derived matrix properties.

The presented methodology is systematic enough to be executed mechanically, that is, without any human intervention, thus setting the ground for a system that indeed automates its application.

Having made clear what is accomplished with this thesis, it should also be pointed out what is beyond its scope: The input to the derivation process is a representation of an iterative method that still has to be derived by an expert in the field. In case of direct methods, the formal description that is input immediately follows from the operation itself, as shown above for the lower triangular system. In contrast, the way from a linear system $Ax = b$ to a formal description of an iterative method that is suitable as input for the presented approach is much more complicated.

1.4. Outline of the Thesis

The thesis is structured as follows: Chapter 2 serves as an introduction to familiarize the reader with the FLAME approach for deriving algorithms for direct methods. In Chapter 3, we lay the foundations for applying a similar approach to iterative methods. The matrix representation for iterative methods is introduced in Section 3.1. In Section 3.2, we describe a method for systematically deriving matrix properties from this representation. At the end of this chapter, in Section 3.3, the implications of the matrix representation for the derivation of algorithms are discussed. In Chapter 4, we explain how to derive algorithms for iterative methods. Finally, Chapter 5 summarizes the results of this thesis and points out opportunities for future research.

2. Derivation of Algorithms for Direct Methods

The purpose of this chapter is to introduce the reader to the systematic derivation of loop-based, blocked algorithms for direct methods using the FLAME methodology. For the derivation, we will follow the systematic approach and the notation of [6]. We begin with the simple example of a triangular linear system in Section 2.1. In Section 2.2, we proceed with a somewhat more elaborate example to go into some details that are not covered in the first one. Next, in Section 2.3, a case will be demonstrated where it is not possible to derive algorithms that compute the solution for every input of the operation. Finally, we discuss the equivalence of loop invariants that are obtained with this approach in Section 2.4.

Notation

Throughout this thesis, two different notations are used for the indexing of matrix blocks and elements. The first one is the standard notation that uses numerals. The second one uses capital letters, as shown below:

$$\begin{array}{lll}
 \bullet (A_L \mid A_R) & \bullet \left(\frac{A_T}{A_B} \right) & \bullet \left(\frac{A_{TL} \mid A_{TR}}{A_{BL} \mid A_{BR}} \right) \\
 \bullet (A_L \mid A_M \mid A_R) & \bullet \left(\frac{A_T}{A_M} \right) & \bullet \left(\frac{A_{TL} \mid A_{TM} \mid A_{TR}}{A_{ML} \mid A_{MM} \mid A_{MR}} \right) \\
 & \bullet \left(\frac{A_T}{A_B} \right) & \bullet \left(\frac{A_{TL} \mid A_{TM} \mid A_{TR}}{A_{BL} \mid A_{BM} \mid A_{BR}} \right)
 \end{array}$$

The subscript letters T, B, L, M and R stand for Top, Bottom, Left, Middle and Right, respectively.

2.1. Triangular Linear System

We begin with deriving algorithms for the linear system $Ax = b$ where A is triangular. The starting point of the derivation is a complete description of the problem, consisting of two logical predicates, a precondition (P_{pre}) and a postcondition (P_{post}). The precondition lists the properties of all quantities that are part of the operation, while the postcondition consists of the equation (or equations) that constitute the operation.

2. Derivation of Algorithms for Direct Methods

Algorithm: ...
{ P _{pre} }
Partition
{ P _{inv} }
While G do
{ (P _{inv}) ∧ (G) }
Repartition
{ P _{before} }
Update
{ P _{after} }
Continue with
{ P _{inv} }
endwhile
{ (P _{inv}) ∧ ¬ (G) }
{ P _{post} }

Figure 2.1.: The skeleton of a FLAME worksheet.

The precondition is true before the execution of the algorithm, and our goal is to find an algorithm that makes the postcondition true upon termination. Additionally, a function is introduced which abstracts from the details of the operation, and instead describes the output as a function of the input. For reasons that will become apparent later, we use \hat{b} to denote the initial content of b . The description for a linear system with A lower triangular is:

$$x := \Phi(A, \hat{b}) \equiv \left\{ \begin{array}{l} P_{\text{pre}} : \{ \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\ \quad \text{LowerTriangular}[A] \wedge \\ \quad \text{Input}[\hat{b}] \wedge \text{Vector}[\hat{b}] \wedge \\ \quad \text{Output}[x] \wedge \text{Vector}[x] \} \\ P_{\text{post}} : \{ Ax = \hat{b} \} \end{array} \right.$$

The actual algorithms are then constructed by filling out a so-called “worksheet”, a template for a loop-based algorithm [2, 6], shown in Figure 2.1.

PME Generation

As a first step towards an algorithm, the PME for the operation is generated. In order to do so, we partition the postcondition. For direct methods, partitioning operands along each dimension in at most two parts proved to be sufficient (this is further discussed in Section 2.4). Since A is lower triangular, and to preserve the triangularity of the resulting objects, a 2×2 partitioning is chosen for it where A_{TL} (and A_{BR}) is

square.

$$\left(\begin{array}{c|c} A_{TL} & 0 \\ \hline A_{BL} & A_{BR} \end{array} \right) \begin{pmatrix} x_T \\ x_B \end{pmatrix} = \begin{pmatrix} \hat{b}_T \\ \hat{b}_B \end{pmatrix}$$

Next, in the so called *Matrix Arithmetic* step, we compute the symbolic multiplication and distribute the equality over the partitioned objects, resulting in the following expression:

$$\left(\begin{array}{c} A_{TL}x_T = \hat{b}_T \\ \hline A_{BL}x_T + A_{BR}x_B = \hat{b}_B \end{array} \right)$$

In the final *Pattern Matching* step, we try to find ways to solve the given equations by writing them as known operations. We begin with the top part. A_{TL} and \hat{b}_T are known, while x_T is unknown. Since A_{TL} is lower triangular, we recognize that $A_{TL}x_T = \hat{b}_T$ is a lower triangular system. Thus, we can rewrite this expression as $x_T := \Phi(A_{TL}, \hat{b}_T)$, utilizing the function introduced earlier. Doing so, we can from now on consider x_T as known, which is necessary to solve the second equation. $A_{BL}x_T$ is a matrix-vector product of two known quantities, and if we subtract it on both sides of the equation, we identify $A_{BR}x_B = \hat{b}_B - A_{BL}x_T$ as another triangular system. Rewriting this as $x_B := \Phi(A_{BR}, \hat{b}_B - A_{BL}x_T)$ provides us with the PME

$$\left(\begin{array}{c} x_T := \Phi(A_{TL}, \hat{b}_T) \\ \hline x_B := \Phi(A_{BR}, \hat{b}_B - A_{BL}x_T) \end{array} \right).$$

Loop Invariant Identification

The next step consists of finding loop invariants (P_{inv}). A loop invariant is a logical predicate that is true at certain points of a loop in an algorithm. It is true before the loop is entered and after it is left, as well as at the beginning and the end of the loop body. It allows to formally reason about the correctness of the loop [8].

As a first step, the PME is decomposed into basic operations, for example a matrix multiplication, or the function representing the operation we are deriving algorithms for. In general, one can think of those basic operations in terms of BLAS-like functions. For the lower triangular system, the following three operations are obtained:

1. $x_T := \Phi(A_{TL}, \hat{b}_T)$
2. $b_B := \hat{b}_B - A_{BL}x_T$
3. $x_B := \Phi(A_{BR}, b_B)$

Then, from those operations, a dependency graph is constructed. Every node represents one operation, while the edges of this directed graph represent the data dependencies between those operations. If one operation requires the output of another operation, the former depends on the latter. Thus, in the graph, there are edges from

2. Derivation of Algorithms for Direct Methods

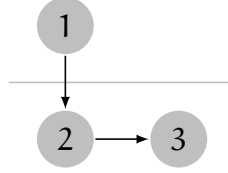


Figure 2.2.: Dependency graph for a lower triangular system.

the node that computes a quantity to those which need this quantity as input. Consider Figure 2.2 for the resulting dependency graph.

Subsets of nodes from this graph are now selected as loop invariants. For reasons that will be explained later, those subsets can neither be empty, nor can they contain all nodes. Furthermore, if a node is part of the subset, then all preceding nodes are also in the set. It follows that there are two subsets of nodes of the graph, $\{1\}$ and $\{1, 2\}$, that can be used as loop invariants.

$$\begin{aligned} \{1\} : P_{\text{inv}}^1 &= \{x_T := \Phi(A_{TL}, \hat{b}_T)\} \\ \{1, 2\} : P_{\text{inv}}^2 &= \{x_T := \Phi(A_{TL}, \hat{b}_T) \wedge b_B := \hat{b}_B - A_{BL}x_T\} \end{aligned}$$

Now, we can also identify the loop guard G . The loop is supposed to terminate when all operands are traversed and the complete solution is computed. Thus, the loop guard depends on how the operands are traversed. The first assignment of the dependency graph operates on x_T , A_{TL} and b_T , so we can conclude that x and b are traversed from the top to the bottom, and A is traversed from the top-left to the bottom-right corner. In the algorithm, this has the result that initially, A_{TL} , x_T and b_T are empty. A_{TL} grows with every iteration until it has the same size as A . Similarly, x_T and b_T grow to the size of x and b . Hence, one possible loop guard is “ $\text{size}(A_{TL}) < \text{size}(A)$ ”. How A_{TL} , x_T and b_T change their sizes is explained in further detail in the next section.

Algorithm Construction

Having found a number of loop invariants, we can proceed to the final step, the derivation of the updates, using the worksheet (Figure 2.1). For each loop invariant, a separate worksheet is filled out. Every loop invariant is rewritten in two different ways, resulting in two predicates per loop invariant. The first predicate specifies the state of all operands before the update (P_{before}). Analogously, the second predicate P_{after} represents the situation after the update. The update is then found by identifying the operations that transform the predicate P_{before} into P_{after} . Intuitively speaking, the difference between both predicates is determined.

The rewritings that are applied to the loop invariant are repartitionings of the already partitioned operands. They are chosen in a way that ensures that the operation makes progress and eventually terminates in combination with a suitable loop guard: Those blocks of the output that are already computed grow, and those parts that are

not computed yet shrink. This is done by splitting off parts of some quantities by applying the “Repartition” rules and merging them with others with the “Continue with” repartitioning.

We begin with repartitioning the partitioned operands. To obtain a blocked algorithm, we use the following “Repartition” rules:

$$\begin{aligned} \left(\begin{array}{c|c} A_{TL} & 0 \\ \hline A_{BL} & A_{BR} \end{array} \right) &\rightarrow \left(\begin{array}{c|c|c} A_{00} & 0 & 0 \\ \hline A_{10} & A_{11} & 0 \\ \hline A_{20} & A_{21} & A_{22} \end{array} \right) \\ \left(\begin{array}{c} x_T \\ \hline x_B \end{array} \right) &\rightarrow \left(\begin{array}{c} x_0 \\ \hline x_1 \\ \hline x_2 \end{array} \right) \\ \left(\begin{array}{c} b_T \\ \hline b_B \end{array} \right) &\rightarrow \left(\begin{array}{c} b_0 \\ \hline b_1 \\ \hline b_2 \end{array} \right) \end{aligned}$$

Doing so, during every iteration, some parts of A_{BR} are split off. To continue after the update, those parts are merged into A_{TL} . Thus, A_{BR} shrinks and A_{TL} grows until the entire matrix, and at the same time both vectors b and x , are traversed. This “Continue with” partitioning is shown below.

$$\begin{aligned} \left(\begin{array}{c|c} A_{TL} & 0 \\ \hline A_{BL} & A_{BR} \end{array} \right) &\leftarrow \left(\begin{array}{c|c|c} A_{00} & 0 & 0 \\ \hline A_{10} & A_{11} & 0 \\ \hline A_{20} & A_{21} & A_{22} \end{array} \right) \\ \left(\begin{array}{c} x_T \\ \hline x_B \end{array} \right) &\leftarrow \left(\begin{array}{c} x_0 \\ \hline x_1 \\ \hline x_2 \end{array} \right) \\ \left(\begin{array}{c} b_T \\ \hline b_B \end{array} \right) &\leftarrow \left(\begin{array}{c} b_0 \\ \hline b_1 \\ \hline b_2 \end{array} \right) \end{aligned}$$

To get to the predicates P_{before} and P_{after} , we have to rewrite the loop invariant using the repartitioned operands. We demonstrate this for the second loop invariant, P_{inv}^2 . For P_{before} , we replace all quantities by their counterparts according to the first partitioning (from 2×2 to 3×3), resulting in

$$\begin{aligned} x_0 &:= \Phi(A_{00}, \hat{b}_0) \\ \left(\begin{array}{c} b_1 \\ \hline b_2 \end{array} \right) &:= \left(\begin{array}{c} \hat{b}_1 \\ \hline \hat{b}_2 \end{array} \right) - \left(\begin{array}{c} A_{10} \\ \hline A_{20} \end{array} \right) x_0. \end{aligned}$$

2. Derivation of Algorithms for Direct Methods

Then, we flatten the expression, that is, perform all algebraic operations, distribute the assignments and decompose the result into its parts. This yields the following predicate:

$$\begin{aligned} P_{\text{before}} = \{ & x_0 := \Phi(A_{00}, \hat{b}_0) \wedge \\ & b_1 := \hat{b}_1 - A_{10}x_0 \wedge \\ & b_2 := \hat{b}_2 - A_{20}x_0 \} \end{aligned}$$

The same is done using the “Continue with” partitioning. To repartition the function Φ , the PME is used.

$$\begin{aligned} \left(\begin{array}{c} x_0 \\ x_1 \end{array} \right) &:= \Phi \left(\left(\begin{array}{c|c} A_{00} & 0 \\ A_{10} & A_{11} \end{array} \right), \left(\begin{array}{c} \hat{b}_0 \\ \hat{b}_1 \end{array} \right) \right) \\ b_2 &:= \hat{b}_2 - (A_{20} \mid A_{21}) \left(\begin{array}{c} x_0 \\ x_1 \end{array} \right) \end{aligned}$$

Thus, we obtain the predicate below:

$$\begin{aligned} P_{\text{after}} = \{ & x_0 := \Phi(A_{00}, \hat{b}_0) \wedge \\ & x_1 := \Phi(A_{11}, \hat{b}_1 - A_{10}x_0) \wedge \\ & b_2 := \hat{b}_2 - A_{20}x_0 - A_{21}x_1 \} \end{aligned}$$

Now, by comparing the predicates P_{before} and P_{after} , we determine the update. Highlighted in red are those parts of P_{after} that do not appear in P_{before} .

$$\begin{aligned} P_{\text{after}} = \{ & x_0 := \Phi(A_{00}, \hat{b}_0) \wedge \\ & \color{red}{x_1 := \Phi(A_{11}, \hat{b}_1 - A_{10}x_0)} \wedge \\ & b_2 := \hat{b}_2 - A_{20}x_0 - \color{red}{A_{21}x_1} \} \end{aligned}$$

This provides us with the following update:

$$\begin{aligned} x_1 &:= \Phi(A_{11}, b_1) \\ b_2 &:= b_2 - A_{21}x_1 \end{aligned}$$

We have now derived a complete algorithm. Consider Figure 2.3 for the filled out worksheet with updates for both loop invariants (omitting some of the repartitionings and all predicates in the interest of visual clarity).

Using the loop guard and the loop invariants, it is easy to see that the derived algorithms are correct and at the end of the computation, the linear system is solved. The algorithms terminate when the loop guard becomes false. The negation of “ $\text{size}(A_{\text{TL}}) < \text{size}(A)$ ” implies that A_{TL} has the same size as A , as it can not be larger. Since A_{TL} is a part of A , this means that A_{TL} is equal to A . Similarly, x_{T} equals x . \hat{b}_{T} and b_{T} are equal

Algorithm: $x := \Phi(A, b)$	
Partition $A \rightarrow \left(\begin{array}{c c} A_{TL} & 0 \\ \hline A_{BL} & A_{BR} \end{array} \right)$ where A_{TL} is 0×0 While $\text{size}(A_{TL}) < \text{size}(A)$ do	
Repartition $\left(\begin{array}{c c} A_{TL} & 0 \\ \hline A_{BL} & A_{BR} \end{array} \right) \rightarrow \left(\begin{array}{c c c} A_{00} & 0 & 0 \\ \hline A_{10} & A_{11} & 0 \\ \hline A_{20} & A_{21} & A_{22} \end{array} \right)$ where A is $k \times k$	
Variant 1	Variant 2
$b_1 := \hat{b}_1 - A_{10}x_0$	$x_1 := \Phi(A_{11}, b_1)$
$x_1 := \Phi(A_{11}, b_1)$	$b_2 := b_2 - A_{21}x_1$
Continue with $\left(\begin{array}{c c} A_{TL} & 0 \\ \hline A_{BL} & A_{BR} \end{array} \right) \leftarrow \left(\begin{array}{c c c} A_{00} & 0 & 0 \\ \hline A_{10} & A_{11} & 0 \\ \hline A_{20} & A_{21} & A_{22} \end{array} \right)$	
endwhile	

Figure 2.3.: Worksheet for a lower triangular linear system.

to \hat{b} and b , respectively. As we updated b , we have to distinguish between its initial and its current content to reason about the correctness of the derived algorithm. All other parts have either size $0 \times n$, 0×0 or $n \times 0$, so they disappear. Plugging that in the loop invariants, in both cases we only get $x := \Phi(A, \hat{b})$, which proves that both algorithms compute the solution to a triangular system.

2.2. Symmetric Positive Definite Linear System

As a second example, we derive algorithms for a symmetric positive definite (SPD) linear system to show in greater detail how the feasibility of loop invariants is checked. Again, the starting point is the formal description of the operation, which is shown

2. Derivation of Algorithms for Direct Methods

below:

$$x := \Sigma(A, b) \equiv \begin{cases} P_{\text{pre}} : \{\text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{SPD}[A] \wedge \\ \quad \text{Input}[b] \wedge \text{Vector}[b] \wedge \\ \quad \text{Output}[x] \wedge \text{Vector}[x]\} \\ P_{\text{post}} : \{\hat{A}x = \hat{b}\} \end{cases}$$

Because of the symmetry of \hat{A} , we apply a 2×2 partitioning where \hat{A}_{TL} (and \hat{A}_{BR}) is square. Furthermore, because of the symmetry, it holds that $\hat{A}_{\text{BL}} = \hat{A}_{\text{TR}}^T$. The following partitioned postcondition is obtained.

$$\left(\frac{\hat{A}_{\text{TL}}x_{\text{T}} + \hat{A}_{\text{TR}}x_{\text{B}} = \hat{b}_{\text{T}}}{\hat{A}_{\text{BL}}x_{\text{T}} + \hat{A}_{\text{BR}}x_{\text{B}} = \hat{b}_{\text{B}}} \right)$$

Unfortunately, none of the two equations matches the function $x := \Sigma(A, b)$. If we rewrite $\hat{A}_{\text{TL}}x_{\text{T}} + \hat{A}_{\text{TR}}x_{\text{B}} = \hat{b}_{\text{T}}$ as $\hat{A}_{\text{TL}}x_{\text{T}} = \hat{b}_{\text{T}} - \hat{A}_{\text{TR}}x_{\text{B}}$, the right-hand side of this equation is not completely known, as x_{B} is not known. For the same reason, the second equation does not match the function either. In order to proceed, we need to perform a number of steps that are not part of the systematic approach presented in [6]. As \hat{A} is SPD and \hat{A}_{TL} is square, we know that \hat{A}_{TL} (and \hat{A}_{BR}) is also SPD, so we can rewrite the top part of the partitioned postcondition as:

$$x_{\text{T}} = \hat{A}_{\text{TL}}^{-1} (\hat{b}_{\text{T}} - \hat{A}_{\text{TR}}x_{\text{B}})$$

By replacing x_{T} in the second equation with the right-hand side of the equation above and performing further manipulations, it is possible to obtain an equation that matches the pattern of a SPD linear system. Since \hat{A} is SPD, we can infer that $\hat{A}_{\text{BR}} - \hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}\hat{A}_{\text{TR}}$ is SPD too [12].

$$(\hat{A}_{\text{BR}} - \hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}\hat{A}_{\text{TR}})x_{\text{B}} = \hat{b}_{\text{B}} - \hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}\hat{b}_{\text{T}}$$

Thus, the following PME is derived:

$$\left(\frac{x_{\text{T}} := \Sigma(\hat{A}_{\text{TL}}, \hat{b}_{\text{T}} - \hat{A}_{\text{TR}}x_{\text{B}})}{x_{\text{B}} := \Sigma(\hat{A}_{\text{BR}} - \hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}\hat{A}_{\text{TR}}, \hat{b}_{\text{B}} - \hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}\hat{b}_{\text{T}})} \right)$$

Libraries for linear algebra operations like BLAS usually do not offer operations for products of more than two quantities. Thus, to decompose this PME into its basic buildings blocks, we need to introduce auxiliary variables. In general, there are multiple ways to compute expressions like $\hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}\hat{A}_{\text{TR}}$, for example by solving $\hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}$ or $\hat{A}_{\text{TL}}^{-1}\hat{A}_{\text{TR}}$ first. To ensure that the dependency graph is as general as possible and does not impose an ordering on expressions like these, multiple auxiliary variables would have to be introduced. Furthermore, their sizes would have to be left unspecified. To keep this example simple, we compute $\hat{A}_{\text{BL}}\hat{A}_{\text{TL}}^{-1}$ first, because it appears twice

2.2. Symmetric Positive Definite Linear System

in the PME. While in that case, we only need one auxiliary variable that has the same size as \hat{A}_{BL} , it is useful to formally introduce a complete matrix Z and partition it in the same way we partition \hat{A} . Doing so, we can treat the auxiliary variable just like all the other operands. Using Z_{BL} , we obtain the following decomposition. The corresponding dependency graph is shown in Figure 2.4.

1. $Z_{BL} := \hat{A}_{BL} \hat{A}_{TL}^{-1}$
2. $A_{BR} := \hat{A}_{BR} - Z_{BL} \hat{A}_{TR}$
3. $b_B := \hat{b}_B - Z_{BL} \hat{b}_T$
4. $x_B := \Sigma(A_{BR}, b_B)$
5. $b_T := \hat{b}_T - \hat{A}_{TR} x_B$
6. $x_T := \Sigma(\hat{A}_{TL}, b_T)$

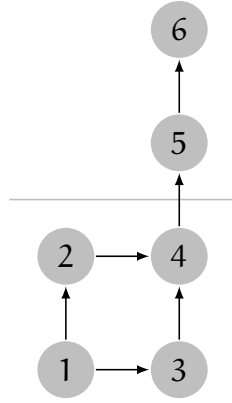


Figure 2.4.: Dependency graph for a SPD system.

The next step is again to select subsets of the dependency graph to use them as loop invariants and determine the loop guard.

Here, \hat{A} is traversed from the bottom right to the top left, and both x and \hat{b} are traversed from the bottom to the top. Thus, the loop guard G is “ $\text{size}(A_{BR}) < \text{size}(A)$ ”. If we had solved $\hat{A}_{BL}x_T + \hat{A}_{BR}x_B = \hat{b}_B$ to x_B earlier, as opposed to x_T , we would have obtained an algorithm that proceeds in the opposite direction.

For this operation, to select loop invariants, the simplified rule that the subsets can neither be empty nor contain all nodes is not sufficient anymore. We now have to consider the full constraints. In general, it has to be checked whether a loop invariant is feasible, that is, if it leads to an algorithm that actually computes the operation. In [6], the author lists the following two constraints:

2. Derivation of Algorithms for Direct Methods

1. There must exist a basic initialization of the operands, i.e., an initial partitioning, that renders the predicate P_{inv} true:

$$\{P_{pre}\}$$

Partition

$$\{P_{inv}\}$$

2. P_{inv} and the negation of the loop guard, G , must imply the postcondition, P_{post} :

$$P_{inv} \wedge \neg G \Rightarrow P_{post}$$

The empty subset always fails to satisfy the second constraint, as it translates to an empty predicate. The empty predicate in conjunction with the negation of the loop guard can never imply the postcondition. Similarly, the set that contains all nodes can not satisfy the first condition. In this case, the solution to the operation would already be computed even before the loop is entered. However, merely partitioning the operands does not render such a loop invariant true. In addition to that, predicates P_{before} and P_{after} would be identical, so no update would be derived.

For SPD linear systems, there are additional subsets that fail to satisfy the second constraint. Let us look at $\{1\}$ for example. The corresponding loop invariant is $P_{inv} = \{Z_{BL} := \hat{A}_{BL}\hat{A}_{TL}^{-1}\}$. The negation of the the loop guard “ $\text{size}(A_{BR}) < \text{size}(A)$ ” implies that A_{BR} is equal to A . Then, \hat{A}_{BL} and \hat{A}_{TL} are empty, and the empty predicate does not imply the postcondition. If we take the subset $\{1, 2\}$, the loop invariant becomes $\{A := \hat{A}\}$ at the end of the operation, which does not imply the postcondition either.

Doing this for all subsets, two feasible loop invariants are obtained:

1. $\{1, 2, 3, 4\} (P_{inv}^1)$
2. $\{1, 2, 3, 4, 5\} (P_{inv}^2)$

As mentioned before, the algorithms proceed from the bottom right to the top left, so the repartitionings are different compared to the ones used for the lower triangular system. First, parts of A_{TL} are split off, which are then merged with A_{BR} after the update. Thus, the “Repartition” statement for A (and Z) is the following:

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array} \right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array} \right)$$

This then is the “Continue with” partitioning:

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array} \right) \leftarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array} \right)$$

Similarly, the repartitionings for b are:

$$\begin{pmatrix} \frac{b_T}{b_B} \end{pmatrix} \rightarrow \begin{pmatrix} \frac{b_0}{b_1} \\ \frac{b_1}{b_2} \end{pmatrix}$$

$$\begin{pmatrix} \frac{b_T}{b_B} \end{pmatrix} \leftarrow \begin{pmatrix} \frac{b_0}{b_1} \\ \frac{b_1}{b_2} \end{pmatrix}$$

x is repartitioned in the same way. In the interest of brevity, the derivation of the updates will not be shown here. There are, however, two important points that should be mentioned. When deriving the predicate P_{before} , the expression

$$\left(\begin{array}{c|c} Z_{20} & Z_{21} \end{array} \right) := \left(\begin{array}{c|c} A_{20} & A_{21} \end{array} \right) \left(\begin{array}{c|c} A_{00} & A_{01} \\ \hline A_{10} & A_{11} \end{array} \right)^{-1}$$

appears. To flatten it, the PME of a different SPD system, namely $XA = B$, has to be used. Then, in order to find out which expressions appear both in P_{before} and P_{after} and to identify the differences, it is necessary to rewrite the expressions of both predicates first. Doing so, the auxiliary variables are eliminated.

Finally, the following two updates are found:

$$\begin{array}{ll} P_{\text{inv}}^1 : b_1 := \hat{b}_1 - A_{12}x_2 & P_{\text{inv}}^2 : x_1 := \Sigma(A_{11}, b_1) \\ x_1 := \Sigma(A_{11}, b_1) & b_0 := b_0 - A_{01}x_1 \end{array}$$

2.3. General Linear System

In this section, we apply the FLAME methodology to a linear system where A is a general, nonsingular matrix.

$$x := \Lambda(A, b) \equiv \begin{cases} P_{\text{pre}} : \{ \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\ \quad \text{Input}[b] \wedge \text{Vector}[b] \wedge \\ \quad \text{Output}[x] \wedge \text{Vector}[x] \} \\ P_{\text{post}} : \{ \hat{A}x = \hat{b} \} \end{cases}$$

Since \hat{A} is not lower triangular, there is more than one possibility for the initial partitioning. The different partitionings are shown in Table 2.1.

At this point, we are again not able to make any further progress with the usual approach. No (sub)expression of the three expressions matches the initial operation. In case of the first two, the parts of \hat{A} are not square. Some or all of the four parts of \hat{A}

2. Derivation of Algorithms for Direct Methods

#	Partitioned postcondition	Flattened expression
1	$(\hat{A}_L \mid \hat{A}_R) \begin{pmatrix} x_T \\ x_B \end{pmatrix} = \hat{b}$	$\hat{A}_L x_T + \hat{A}_R x_B = \hat{b}$
2	$\begin{pmatrix} \hat{A}_T \\ \hat{A}_B \end{pmatrix} x = \begin{pmatrix} \hat{b}_T \\ \hat{b}_B \end{pmatrix}$	$\begin{pmatrix} \hat{A}_T x = \hat{b}_T \\ \hat{A}_B x = \hat{b}_B \end{pmatrix}$
3	$\begin{pmatrix} \hat{A}_{TL} & \hat{A}_{TR} \\ \hat{A}_{BL} & \hat{A}_{BR} \end{pmatrix} \begin{pmatrix} x_T \\ x_B \end{pmatrix} = \begin{pmatrix} \hat{b}_T \\ \hat{b}_B \end{pmatrix}$	$\begin{pmatrix} \hat{A}_{TL} x_T + \hat{A}_{TR} x_B = \hat{b}_T \\ \hat{A}_{BL} x_T + \hat{A}_{BR} x_B = \hat{b}_B \end{pmatrix}$

Table 2.1.: Possible partitionings for a general linear system.

in the third expression could be square, but since the partitioning does not prescribe any sizes, we can in general not assume that this is the case for any part.

If we use the third partitioning and assume that \hat{A}_{TL} (and \hat{A}_{BR}) is square, we are in a similar situation as with the SPD system. Recall that to proceed, we had to rewrite

$$\hat{A}_{TL} x_T + \hat{A}_{TR} x_B = \hat{b}_T$$

as

$$x_T = \hat{A}_{TL}^{-1} (\hat{b}_T - \hat{A}_{TR} x_B).$$

This was possible because \hat{A} is SPD, so we were able to infer that \hat{A}_{TL} is SPD, and thus nonsingular, as well. The difference now is that the \hat{A}_{TL} part of a nonsingular matrix \hat{A} is in general not nonsingular. While we could assume that it is, proceed as we did in the previous section and derive similar algorithms, those algorithms would not solve general linear systems.

2.4. Equivalence of Loop Invariants

For direct methods, one observes that finer partitionings lead to additional loop invariants. A finer partitioning here means that we partition along one dimension in more than the usual two parts. Some of those loop invariants lead to truly new algorithms that cannot be derived with a coarser partitioning. Intuitively, one can say that those algorithms expose intermediate steps that are not exposed with a coarser partitioning. Others are in some sense redundant, because they lead to algorithms that are computationally equivalent to algorithms derived from different loop invariants. We then consider those loop invariants to be equivalent. Those loop invariants may or may not have the same granularity. This means that two loop invariants from a 3×3 partitioning can be equivalent, but it is also possible that a 2×2 loop invariant is equivalent to a loop invariant obtained from a 3×3 partitioning.

This is relevant for the derivation of algorithms for iterative methods because we compare different partitionings in Section 3.3. Even more important, the partitioning that is used for iterative methods behaves differently in this regard, as discussed in Section 4.3.1.

In this section, we show what equivalence means in this context and under which conditions this behavior occurs. We begin by introducing the notion of equivalence for algorithms and then extend the results to loop invariants.

Equivalence of Algorithms

The triangular continuous-time Sylvester equation, which will serve as an example, is defined as $AX + XB = C$ with A and B being triangular. In the following, we assume A and B to be lower and upper triangular, respectively:

$$X := \Psi(A, B, C) \equiv \begin{cases} P_{\text{pre}} : \{ \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{LowerTriangular}[A] \wedge \\ \quad \text{Input}[B] \wedge \text{Matrix}[B] \wedge \text{UpperTriangular}[B] \wedge \\ \quad \text{Input}[C] \wedge \text{Matrix}[C] \wedge \\ \quad \text{Output}[X] \wedge \text{Matrix}[X] \} \\ P_{\text{post}} : \{ AX + XB = \hat{C} \} \end{cases}$$

Applying a 1×3 partitioning to X results in the following PME:

$$\left(\begin{array}{c|c|c} X_L = \Psi(A, B_{TL}, \hat{C}_L) & X_M = \Psi(A, B_{MM}, \hat{C}_M - X_L B_{TM}) & X_R = \Psi(A, B_{BL}, \hat{C}_R - X_L B_{TR} - X_M B_{MR}) \end{array} \right)$$

At this point, we introduce a new notation for loop invariants taken from [6]. Instead of a logical predicate, the PME is used, leaving out those (sub)expressions that are not part of the loop invariant. If entire blocks are not included, we use the symbol \neq to express that no constraints are imposed on this part. We now choose the following two loop invariants and derive the corresponding updates, skipping the intermediate steps.

$$\begin{aligned} P_{\text{inv}} &= (X_L = \Psi(A, B_{TL}, \hat{C}_L) \mid \neq \mid \neq) \\ P'_{\text{inv}} &= (X_L = \Psi(A, B_{TL}, \hat{C}_L) \mid X_M = \Psi(A, B_{MM}, \hat{C}_M - X_L B_{TM}) \mid \neq) \end{aligned}$$

The updates are for P_{inv} and P'_{inv} are:

$$\begin{aligned} P_{\text{inv}} : C_1 &:= \hat{C}_1 - X_0 B_{01} \\ X_1 &:= \Psi(A, B_{11}, C_1) \\ P'_{\text{inv}} : C_2 &:= \hat{C}_2 - X_0 B_{02} - X_1 B_{12} \\ X_2 &:= \Psi(A, B_{22}, C_2) \end{aligned}$$

2. Derivation of Algorithms for Direct Methods

Obviously, the updates for X_1 and X_2 only differ in the indices, which are shifted by one. B_{22} is the next block following B_{11} on the main diagonal of B , and C_2 is the next set of columns following C_1 in C . Considering that those updates happen inside a loop, as the computation unfolds, they will both point to the same parts of B and C , respectively.

Written as above, the updates for C do not immediately appear equivalent. Using BLAS operations, C_1 would be updated with one call to GEMM, C_2 with two. However, when we rewrite the update for C_2 as follows, it is easy to see that it is equivalent to just one GEMM operation.

$$C_2 := \hat{C}_2 - (X_0 \mid X_1) \begin{pmatrix} B_{02} \\ B_{12} \end{pmatrix}$$

Similarly to C_1 and C_2 , X_0 and $(X_0 \mid X_1)$ will, at some point during the computation, contain the same parts of X . The same is true for B_{01} and

$$\begin{pmatrix} B_{02} \\ B_{12} \end{pmatrix}.$$

Hence, we consider both algorithms to be equivalent. Note that in general, the equivalence can be much less obvious, requiring much more elaborate rewriting.

In the example above, it is easy to see that simply by replacing quantities in one update with the corresponding quantities of the other update, it is possible to transform one update into the other. The only constraint is that if one quantity replaces another, both must have matching sizes. Since we do not specify any concrete dimension and the size of some parts changes during the computation, we have to compare them using a symbolic representation of those sizes. Let us begin with listing the dimensions of all operands of the Sylvester equation:

- X and C are of size $n \times m$.
- A is of size $n \times n$.
- B is of size $m \times m$.

After the repartitioning, we specify the sizes as follows.

$$\begin{array}{c} \begin{array}{cccc} b_i & b & b & b_{j+c} \end{array} \\ n \left(X_0 \mid X_1 \mid X_2 \mid X_3 \right) \end{array}$$

$$\begin{array}{c} \begin{array}{cccc} b_i & b & b & b_{j+c} \end{array} \\ n \left(C_0 \mid C_1 \mid C_2 \mid C_3 \right) \end{array}$$

$$\begin{array}{c}
 \begin{array}{cccc}
 & bi & b & b & bj+c \\
 bi & \left(\begin{array}{c|c|c|c} B_{00} & B_{01} & B_{02} & B_{03} \\ \hline 0 & B_{11} & B_{12} & B_{13} \\ \hline 0 & 0 & B_{22} & B_{23} \\ \hline 0 & 0 & 0 & B_{33} \end{array} \right) \\
 b \\
 b \\
 bj+c
 \end{array}
 \end{array}$$

b is the block size, bi and bj are unspecified multiples of the block size. c is an additional constant that is nonzero if the block size b does not divide m and/or n . Blocks with constant sizes, that is, any combination of n , m and b , can only be replaced with blocks of the corresponding quantities that have the exact same sizes. For example, it is possible to replace X_1 with X_2 , or B_{22} with B_{11} .

Similarly, the dimensions have to match for quantities with variable sizes. If one dimension is a multiple of the block size, any other multiple of the block size matches. As an example, bi matches $bi + b$, and $bj + c + 2b$ matches $bj + c$. It is important to note that bi does not match b , because the former is variable, and the latter is constant. Thus, two valid replacements are:

$$\begin{aligned}
 B_{00} &\rightarrow \left(\begin{array}{c|c} B_{00} & B_{01} \\ \hline 0 & B_{11} \end{array} \right) \\
 (C_2 \mid C_3) &\rightarrow C_3
 \end{aligned}$$

Returning to the initial example, the replacement to transform the update for the loop invariant P_{inv} into the one for the loop invariant P'_{inv} is the following:

$$\begin{aligned}
 B_{01} &\rightarrow \left(\begin{array}{c} B_{02} \\ \hline B_{12} \end{array} \right) \\
 B_{11} &\rightarrow B_{22} \\
 C_1 &\rightarrow C_2 \\
 X_0 &\rightarrow (X_0 \mid X_1) \\
 X_1 &\rightarrow X_2
 \end{aligned}$$

Equivalence of Loop Invariants

As mentioned earlier, we consider two loop invariants to be equivalent if the resulting algorithms are equivalent. To determine this equivalence, however, it is not necessary to derive algorithms. The same method of replacing quantities can be performed directly on the loop invariants. If, with such a replacement, one loop invariant can be transformed into another one, both would result in the same update. Consequently, the loop invariants themselves are equivalent.

Let us demonstrate this for the loop invariants P_{inv} and P'_{inv} . We choose the follow-

2. Derivation of Algorithms for Direct Methods

ing replacements:

$$\begin{aligned} X_L &\rightarrow (X_L \mid X_M) \\ C_L &\rightarrow (C_L \mid C_M) \\ B_{TL} &\rightarrow \left(\begin{array}{c|c} B_{TL} & B_{TM} \\ \hline 0 & B_{MM} \end{array} \right) \end{aligned}$$

Applying those to the loop invariant

$$P_{\text{inv}} = (X_L = \Psi(A, B_{TL}, \hat{C}_L) \mid \neq \mid \neq),$$

the equation for computing X_L becomes

$$(X_L \mid X_M) = \Psi \left(A, \left(\begin{array}{c|c} B_{TL} & B_{TM} \\ \hline 0 & B_{MM} \end{array} \right), (\hat{C}_L \mid \hat{C}_M) \right).$$

Flattening the expressions, we obtain the second loop invariant, P'_{inv} :

$$P'_{\text{inv}} = (X_L = \Psi(A, B_{TL}, \hat{C}_L) \mid X_M = \Psi(A, B_{MM}, \hat{C}_M - X_L B_{TM}) \mid \neq)$$

Equivalence of Loop Invariants of Different Granularities

The presented method of term rewriting can also be used to decide whether an algorithm or loop invariant is equivalent to a different one obtained with a finer or coarser partitioning. So far, the quantities on both sides of the replacements originated from the same partitioned object. Now, each operand is partitioned twice, using partitionings of two different granularities. Parts obtained from one are replaced with parts of the other. Apart from that, exactly the same rules for the replacement hold.

We demonstrate this with a more involved example: The inverse of a lower triangular matrix. A 2×2 partitioning yields the following PME, using \tilde{X} and \tilde{L} to avoid confusion:

$$\left(\begin{array}{c|c} \tilde{X}_{TL} := \tilde{L}_{TL}^{-1} & 0 \\ \hline \tilde{X}_{BL} := -\tilde{L}_{BR}^{-1} \tilde{L}_{BL} \tilde{X}_{TL} & \tilde{X}_{BR} := \tilde{L}_{BR}^{-1} \end{array} \right) \quad (2.1)$$

This is the 3×3 PME:

$$\left(\begin{array}{c|c|c} X_{TL} := L_{TL}^{-1} & 0 & 0 \\ \hline X_{ML} := -L_{MM}^{-1} L_{ML} X_{TL} & X_{MM} := L_{MM}^{-1} & 0 \\ \hline X_{BL} := -L_{BR}^{-1} L_{BM} X_{ML} - L_{BR}^{-1} L_{BL} X_{TL} & X_{BM} := -L_{BR}^{-1} L_{BM} X_{MM} & X_{BR} := L_{BR}^{-1} \end{array} \right)$$

We will show that the two loop invariants below are equivalent:

2.4. Equivalence of Loop Invariants

$$\left(\begin{array}{c|c} \tilde{X}_{TL} := \tilde{L}_{TL}^{-1} & 0 \\ \hline \tilde{X}_{BL} := -\tilde{L}_{BR}^{-1} \tilde{L}_{BL} \tilde{X}_{TL} & \neq \end{array} \right) \quad \left(\begin{array}{c|c|c} X_{TL} := L_{TL}^{-1} & 0 & 0 \\ \hline X_{ML} := -L_{MM}^{-1} L_{ML} X_{TL} & \neq & 0 \\ \hline X_{BL} := -L_{BR}^{-1} L_{BM} X_{ML} - L_{BR}^{-1} L_{BL} X_{TL} & \neq & \neq \end{array} \right)$$

To do so, we first need to choose an appropriate replacement. The replacement used in this example will transform the 2×2 loop invariant into the 3×3 loop invariant, so some parts of the coarser loop invariant will be replaced with multiple parts of the finer one.

$$\tilde{X}_{TL} \rightarrow X_{TL} \quad (2.2)$$

$$\tilde{L}_{TL} \rightarrow L_{TL} \quad (2.3)$$

$$\tilde{X}_{BL} \rightarrow \begin{pmatrix} X_{ML} \\ X_{BL} \end{pmatrix} \quad (2.4)$$

$$\tilde{L}_{BL} \rightarrow \begin{pmatrix} L_{ML} \\ L_{BL} \end{pmatrix} \quad (2.5)$$

$$\tilde{L}_{BR} \rightarrow \left(\begin{array}{c|c} L_{MM} & 0 \\ \hline L_{BM} & L_{BR} \end{array} \right) \quad (2.6)$$

To make sure that it is a valid replacement, it is necessary to check if the dimensions of the objects on both sides of the replacements match. The sizes of L , and likewise X , are as follows:

$$\begin{array}{cc} & \begin{array}{cc} bi & bj+c \end{array} \\ \begin{array}{c} bi \\ bj+c \end{array} & \left(\begin{array}{c|c} \tilde{L}_{TL} & 0 \\ \hline \tilde{L}_{BL} & \tilde{L}_{BR} \end{array} \right) \end{array} \quad \begin{array}{cc} & \begin{array}{ccc} bi & b & bj+c \end{array} \\ \begin{array}{c} bi \\ b \\ bj+c \end{array} & \left(\begin{array}{c|c|c} L_{TL} & 0 & 0 \\ \hline L_{ML} & L_{MM} & 0 \\ \hline L_{BL} & L_{BM} & L_{BR} \end{array} \right) \end{array}$$

The replacements (2.2) and (2.3) are easily seen to be correct, as all parts have size $bi \times bi$. In case of (2.4) and (2.5), the left-hand sides have the dimensions $(bj + c) \times b$, and the right-hand sides $(bj + c + b) \times b$. The number of rows matches because both are a multiple of the block size b , plus a constant c . Similarly, the last replacement (2.6) is valid, as the sizes $(bj + c) \times (bj + c)$ and $(bj + c + b) \times (bj + c + b)$ are equivalent.

Applying this replacement to the coarser loop invariant, we obtain two assignments:

$$X_{TL} := L_{TL}^{-1} \quad (2.7)$$

$$\begin{pmatrix} X_{ML} \\ X_{BL} \end{pmatrix} := - \left(\begin{array}{c|c} L_{MM} & 0 \\ \hline L_{BM} & L_{BR} \end{array} \right)^{-1} \begin{pmatrix} L_{ML} \\ L_{BL} \end{pmatrix} X_{TL} \quad (2.8)$$

Clearly, the first one already has the exact same shape as in the finer loop invariant. The second one, however, requires some rewriting. To be able to compute the product

2. Derivation of Algorithms for Direct Methods

on the right-hand side, we first have to find the symbolic inverse of

$$\left(\begin{array}{c|c} L_{MM} & 0 \\ \hline L_{BM} & L_{BR} \end{array} \right).$$

We obtain it from the PME of the inverse of a lower triangular matrix, (2.1), by eliminating all occurrences of parts of X and replacing the corresponding quantities:

$$\left(\begin{array}{c|c} L_{MM} & 0 \\ \hline L_{BM} & L_{BR} \end{array} \right)^{-1} = \left(\begin{array}{c|c} L_{MM}^{-1} & 0 \\ \hline -L_{BR}^{-1}L_{BM}L_{MM}^{-1} & L_{BR}^{-1} \end{array} \right)$$

Assignment (2.8) then becomes

$$X_{ML} := -L_{MM}^{-1}L_{ML}X_{TL} \quad (2.9)$$

$$X_{BL} := L_{BR}^{-1}L_{BM}L_{MM}^{-1}L_{ML}X_{TL} - L_{BR}^{-1}L_{BL}X_{TL}. \quad (2.10)$$

The assignment for X_{ML} is now identical to the one in the 3×3 loop invariant. While the assignments for X_{BL} still differ, we observe that it is possible to replace a subexpression of (2.10) with $-X_{ML}$:

$$X_{BL} := L_{BR}^{-1}L_{BM} \underbrace{L_{MM}^{-1}L_{ML}X_{TL}}_{-X_{ML}} - L_{BR}^{-1}L_{BL}X_{TL}$$

Thus, we obtain the following expression, which is the same as the assignment for X_{BL} in the finer loop invariant:

$$X_{BL} := -L_{BR}^{-1}L_{BM}X_{ML} - L_{BR}^{-1}L_{BL}X_{TL}$$

3. Derivation of Algorithms for Iterative Methods: Foundations

To be able to use a FLAME-like methodology to derive algorithms for iterative methods, a matrix representation of those methods is required. Such a representation was used in [5] for CG and the Krylov sequence, and in [4] for some CG variants.

This chapter begins with the introduction of an additional notation. The details of the matrix representation are explained in the first section. In the second part, a systematic method for deriving properties of matrices from this representation is presented. Finally, different possible partitionings are discussed, in addition to their implications for the derivation of algorithms.

Notation

In this thesis, we use a notation that deviates slightly from the one used in [4] and [5]. We use e_0 to denote the unit vector that is one in the first position and e_r for the unit vector that is one in the last position. Both are column vectors. The matrix J is a square matrix with ones on the lower diagonal:

$$J = \begin{pmatrix} 0 & 0 & \dots \\ 1 & 0 & \\ 0 & 1 & \ddots \\ \vdots & 0 & \ddots \end{pmatrix}$$

As usual, I is the identity matrix. The dimension of those matrices and vectors will not be given explicitly if they are clear from the context. If X is a matrix, we use \underline{X} to indicate that the right-most column of this matrix is omitted. Thus, \underline{I} and \underline{J} are both lower trapezoidal matrices with one more row than columns.

3.1. Matrix Representation

In this section, we discuss the details of the matrix representation for iterative methods used in this thesis. Note that the representation for CG, which will serve as an example, is a slightly modified version of the one introduced in [5]. The difference lies

3. Derivation of Algorithms for Iterative Methods: Foundations

in the use of the underline. The three governing equations are shown below. Deriving those equations is not trivial and beyond the scope of this thesis.

$$APD = R (\underline{I} - \underline{J}) \quad (3.1)$$

$$P (I - U) = \underline{R} \quad (3.2)$$

$$PD = X (\underline{I} - \underline{J}) \quad (3.3)$$

The operands have the following properties:

- $A \in \mathbb{R}^{n \times n}$ is the coefficient matrix of the linear system that is supposed to be solved. It is nonsingular. Depending on whether A is symmetric or not, different algorithms can be derived.
- $P \in \mathbb{R}^{n \times m}$ is the matrix of search directions, that is, each column represents the search direction vector during one iteration. It is initially unknown.
- $D \in \mathbb{R}^{m \times m}$ is an unknown diagonal matrix.
- $R \in \mathbb{R}^{n \times (m+1)}$ is the residual matrix. Initially, only the first column r_0 is known. It is computed as $r_0 = Ax_0 - b$, where x_0 is an initial guess for the solution. Additionally, it is orthogonal.
- $U \in \mathbb{R}^{m \times m}$ is unknown and upper diagonal¹ if A is symmetric. Otherwise, it is strictly upper triangular.
- $X \in \mathbb{R}^{n \times (m+1)}$ is the matrix of approximated solution vectors. Similar to R , only the first column is initially known.

While it might seem unusual that the same matrix (R) appears twice in the governing equations with varying sizes, this is necessary to ensure the correctness of the last column of R and X . The formula for computing the residual, in indexed notation, is $r_{i+1} = r_i - Ap_i\delta_i$. Without the additional column of R in equation (3.1), the incorrect equation $Ap_i\delta_i = r_i$ would be obtained for the last iteration. Similarly, from equation (3.3), we would obtain $p_i\delta_i = x_i$, which is not correct either.

One of the fundamental differences compared to direct methods is that the dimensions of some matrices are not fixed. Usually, all dimensions are determined by the sizes of input operands. In case of the LU-factorization, for example, we know that L and U have the same size as A . If not, the equation $LU = A$ is not valid. If A is of size $n \times n$, and the LU factorization of a $k \times k$ block of A , with $k < n$, is computed, then the postcondition is not rendered true.

Due to the orthogonality of R , $m + 1$ can not be larger than n , as the number of n -dimensional, orthogonal vectors is at most n . However, for every $m < n$, the equations above can be satisfied by performing the corresponding number of iterations.

¹This property is defined in Appendix A.

For the systematic derivation of algorithms, this introduces the problem that it is not possible to derive a loop guard exclusively from the equations. It makes no sense to compare the size of a block of R to R itself, because R grows too. For those iterative methods that are used to find solutions for linear systems, the goal is usually to minimize the residual in some norm [1]. Thus, the loop guard typically is a predicate like “ $\|r_i\| \geq \varepsilon$ ”, where ε is a threshold chosen by the user. We make sure that the postcondition of CG correctly represents the situation at the end of the operation by adding $\|Re_r^T\| < \varepsilon$ to it. This also allows us to derive a suitable loop guard from it.

For stationary iterative methods, the stopping criterion quite often is $\|x_i - x_{i-1}\| < \varepsilon$. Translating that in our notation, we obtain $\|Xe_r^T - Xe_{r-1}^T\| < \varepsilon$. For iterative methods where those criteria are not applicable, we will add an expression to the postcondition that fixes the number of columns of a matrix with variable size. Such a predicate could be “ $\text{size}(Y) = n \times k$ ”.²

3.2. Systematic Derivation of Matrix Properties

In order to derive algorithms for CG from its description in matrix form, it is necessary to use properties of matrices and expressions that are not explicitly part of the initial description. Those properties have to be derived from the description by means of algebraic manipulation and deductive reasoning. To automate the process of finding algorithms, an automatic method for the derivation of properties is necessary. Thus, this systematic approach should replicate the steps performed by a human expert, without actually requiring any human guidance. In this section, we describe such a method.

3.2.1. Preliminaries

To begin, it is useful to formalize the notion of properties and equations. We start with the most basic building blocks, terms and expressions.

Definition 3.2.1 (Terms and Expressions)

Terms are inductively defined as follows:

1. Every matrix, vector and scalar is a term.
2. If t is a term, then t^T , $(-t)$ and t^{-1} are terms.
3. If t_1 and t_2 are terms, then $(t_1 + t_2)$ and $(t_1 \cdot t_2)$ are terms.

Every term is also an *expression*. Furthermore, if t_1 and t_2 are terms, then $t_1 = t_2$ is an *expression*. ■

²It is not uncommon to combine multiple criteria [1]. For the sake of simplicity, we do not use more than one at a time.

3. Derivation of Algorithms for Iterative Methods: Foundations

Note that this is a simplified definition, as it does not include functions, nor does it make any statement about the validity of terms. A product of two matrices with dimensions that do not match is still a valid term. Nonetheless, this definition is sufficient for our purposes. Furthermore, we usually use a simplified notation, omitting parentheses, if they are unnecessary, as well as the multiplication dot. We can now define equations and properties:

Definition 3.2.2 (Equations)

Let t_1, t_2 be terms. An *equation* is an expression of the form $t_1 = t_2$. ■

Definition 3.2.3 (Properties)

Let t be a term and P be a boolean predicate. Then $P[t]$ is a *property*. ■

Note that properties are always predicates, even if they can also be expressed as equations. Take the property $\text{Orthogonal}[R]$ as an example. It implies $\text{Diagonal}[R^T R]$. Applying the partitioning that is also applied to the postcondition,

$$R \rightarrow (R_L \mid r_M \mid R_R),$$

to $R^T R$, we find out, among other things, that $R_L^T R_L$ is diagonal and $r_M^T R_L$ equals zero. Instead of considering the equation $r_M^T R_L = 0$ as a property, we use $\text{Zero}[r_M^T R_L]$. This allows us to use a more consistent notation and simplifies the systematic derivation.

3.2.2. Representing Knowledge about Linear Algebra

A human expert who derives properties of matrices inevitably applies some basic knowledge about linear algebra. To allow a system to replicate the expert reasoning, it needs a knowledgebase that encodes this knowledge. We define five different types of implications that will be included in the knowledgebase.

1. $P_1[t] \wedge \dots \wedge P_i[t] \rightarrow P[t]$

This type of implication allows to reason about the combination of properties of one single term. One example is:

$$\text{LowerTriangular}[t] \wedge \text{Symmetric}[t] \rightarrow \text{Diagonal}[t] \quad (3.4)$$

2. $P_1[t_1] \wedge \dots \wedge P_i[t_i] \wedge \exists t \rightarrow P[t]$

Here, t_1, \dots, t_i are subterms of t . Thus, it allows the system to infer the properties of a product or sum of multiple quantities with different properties. The $\exists t$ is used to avoid deriving properties for terms that do not occur anywhere. Consider two examples:

$$\text{Diagonal}[t_1] \wedge \text{LowerTriangular}[t_2] \wedge \exists t_1 t_2 \rightarrow \text{LowerTriangular}[t_1 t_2] \quad (3.5)$$

$$\text{StrictlyUpperTriangular}[t] \wedge \exists I + t \rightarrow \text{UpperTriangular}[I + t] \quad (3.6)$$

3. $P[t_1] \wedge t_1 = t_2 \rightarrow P[t_2]$

This implication enables the system to propagate properties across equalities. P now is a pattern that matches any property.

4. $P_1[t] \rightarrow P_2[f(t)]$

f is a function, for example transposition. Thus, this kind of implication allows to derive properties of transposed or inverted quantities. In addition to that, it is used reason about orthogonal or orthonormal matrices. Let us look at three examples:

$$\text{LowerTriangular}[t] \rightarrow \text{UpperTriangular}[t^T] \quad (3.7)$$

$$\text{LowerTriangular}[t] \rightarrow \text{LowerTriangular}[t^{-1}] \quad (3.8)$$

$$\text{Orthogonal}[t] \rightarrow \text{Diagonal}[t^T t] \quad (3.9)$$

5. $t \rightarrow f(t)$

For one of the steps of the method presented in this chapter, it is necessary that properties have a canonical form: The unary operators $^{-1}$ and T will not be applied to products. However, some of the other types of implications may produce such terms. To transform those terms into the canonical form, implications are necessary that distribute those unary operators across products. Those implications are:

$$(t_1 t_2)^T \rightarrow t_2^T t_1^T \quad (3.10)$$

$$(t_1 t_2)^{-1} \rightarrow t_2^{-1} t_1^{-1} \quad (3.11)$$

Note that all the terms in the implications above may match not only a single operand, but every term that has the required property. So if the product AB is lower triangular, implication (3.7) tells us that $(AB)^T$ is upper triangular, and using (3.10), we find out that $B^T A^T$ is upper triangular.

3.2.3. Derivation of Properties

In the following, we will assume that all equations have the form

$$\prod_{1 \leq i \leq n} t_i = \prod_{1 \leq j \leq m} t'_j.$$

In case of a sum $X + Y$, the product consists of just one term $t = X + Y$. This assumption is no restriction as the only sums that appear in the matrix representations of iterative methods are of the form $I - Y$ and are used exclusively to emphasize the structure of those matrices.

3. Derivation of Algorithms for Iterative Methods: Foundations

Initialization We start with two sets, \mathcal{P} and \mathcal{E} . \mathcal{P} is a subset of the precondition P_{pre} of the description of an operation. It only contains those properties that describe operand types, like $\text{Square}[X]$, $\text{Matrix}[X]$ or $\text{Diagonal}(X)$. Not included are properties that specify what is input and output, as they are superfluous for the derivation. \mathcal{E} is a set of expressions which contains the equations of the corresponding postcondition P_{post} .

Derivation of Properties The implications of the knowledgebase \mathcal{K} are used to derive all possible properties at this stage. This is done as follows:

- Given $P_1[t] \wedge \dots \wedge P_i[t] \rightarrow P[t] \in \mathcal{K}$, the property $P[t]$ is added to \mathcal{P} if $P_k[t] \in \mathcal{P}$ for all $k \in \{1, \dots, i\}$.
- Given $P_1[t] \wedge \dots \wedge P_i[t] \wedge \exists t \rightarrow P[t] \in \mathcal{K}$, the property $P[t]$ is added to \mathcal{P} if
 1. there is an equation $e \in \mathcal{E}$ that contains the term t , and
 2. $P_k[t] \in \mathcal{P}$ for all $k \in \{1, \dots, i\}$.
- Given $P[t_1] \wedge t_1 = t_2 \rightarrow P[t_2] \in \mathcal{K}$, the property $P[t_2]$ is added to \mathcal{P} if $P[t_1] \in \mathcal{P}$ and $(t_1 = t_2) \in \mathcal{E}$.
- Given $P_1[t] \rightarrow P_2[f(t)] \in \mathcal{K}$, the property $P_2[f(t)]$ is added to \mathcal{P} if $P_1[t] \in \mathcal{P}$.
- Given $t \rightarrow f(t) \in \mathcal{K}$ and $P[t] \in \mathcal{P}$, the property $P[f(t)]$ is added to \mathcal{P} .

Matrix Inversion Then, new expressions are added to \mathcal{E} .

1. For every equation $e \in \mathcal{E}$ with $e = (t_1 \cdots t_i = u_1 \cdots u_j)$, it is checked whether t_1 , t_i , u_1 or u_j is invertible.
2. If t_1 or u_1 is invertible, t_1^{-1} or u_1^{-1} , respectively, is multiplied from the left to e , eliminating the invertible term on the one side and adding its inverse on the other. We proceed analogously with t_i and u_j , then multiplying the inverse of the term from the right.
3. The resulting, new equation e' is added to \mathcal{E} .

Those three steps are repeated until no new expressions are found.

Application of Properties In the final step, we apply known properties to expressions to derive new properties. Intuitively, we multiply quantities to both sides of an equation in order to recreate subexpressions that are known to present some property, which are then used to infer new properties.

1. For every property $P[t] \in \mathcal{P}$ with $t = t_1 \cdots t_i$, $i > 1$ a set of tuples

$$S(t) = \{(t_1 \cdots t_k, t_{k+1} \cdots t_i) \mid 1 \leq k < i\}$$

is generated. Intuitively, this set contains the term t , split into two parts in all possible ways.

2. Let $(t_L, t_R) \in S(t)$. If there is an equation $e \in \mathcal{E}$ where t_L is the rightmost (sub)term in any term, then t_R is multiplied from the right. t_L is multiplied from the left in the analogous case. Let e' be the resulting equation. As an example, let e be $ABC = D$ and $(BC, F) \in S(t)$. Since BC appears as the rightmost subterm in e , $t_R = F$ is multiplied from to right, resulting in $ABCF = DF$.
3. The knowledgebase is used to derive new properties using e' that are added to \mathcal{P} (see step "Derivation of Properties"). e' is *not* added to \mathcal{E} .

3.2.4. Design Considerations

One observes that the presented method is not goal-oriented. The derived properties are mainly used to solve equations, so it might seem more natural to derive properties starting with an equation that has to be solved. Based on this equation, expressions would be selected, and in a second step, the properties of those expression would be derived. Those properties would then be used to solve the equation. Instead, the presented method derives a large number of properties, irrespective of the question whether they might be useful or not.

The problem with a goal-oriented approach is that quite often, it is not obvious from the equation which property could be used to solve it. This leaves us with the much more challenging task of identifying which expression might have relevant properties. Take the following equation as an example, which appears when deriving algorithms for nonsymmetric CG. P_L , p_M , r_M and A are known.

$$-P_L u_{TM} + p_M = r_M$$

It is solved for u_{TM} by using the fact that $P^T A P$ is lower triangular. Thus, $P_L^T A P_L$ is lower triangular as well and $P_L^T A p_M$ is zero. Multiplying $P_L^T A$ from the left to both sides of the equation gives us

$$-P_L^T A P_L u_{TM} = P_L^T A r_M.$$

This now is a triangular system that can easily be solved. While it is possible to individually derive that $P_L^T A P_L$ is lower triangular and $P_L^T A p_M$ is zero, the initial equation gives us little to no indication to inspect the properties of those expressions in the first place.

One the other hand, the advantage of a method that is not goal-oriented is that it may find properties that we do not expect to find.

Orthogonality of the Residual Matrix

For some iterative methods, for example CG, the residual matrix R is orthogonal. In the postcondition of those methods, R usually appears multiple times, either as a

3. Derivation of Algorithms for Iterative Methods: Foundations

whole, or without the last column (\underline{R}). Unfortunately, for the derivation of properties, this poses a problem.

If R is orthogonal, then $R^T R$ and $\underline{R}^T \underline{R}$ are diagonal. $R^T \underline{R}$ and its transpose are rectangular and all entries except for the ones on the main diagonal are zero. Thus, in some sense, they are diagonal as well.

It turns out that for deriving certain properties, $R^T \underline{R}$ is needed. Unfortunately, using the described method, neither `Orthogonal[R]` nor `Orthogonal[\underline{R}]` implies any property of $R^T \underline{R}$. One way to solve this would be to treat any matrix that appears with and without the last column in a special way, such that properties of $R^T \underline{R}$ and $\underline{R}^T R$ are found as well. This, however, would require significant modifications of the derivation process.

The simpler solution is to consider R and \underline{R} to be two distinct objects, and properties of $R^T \underline{R}$ and its transpose are added to the precondition.

Substituting Equations

Note that we deliberately avoid substituting quantities in one equation by expressions obtained from others. While this might be a very natural approach if deriving properties by hand, doing this systematically is difficult. Plugging in equations quickly leads to arbitrarily large expressions unless some heuristics are applied to terminate this process. Apart from that, it is possible to achieve the same results using the approach presented in this section. Consider a short example to get an intuition why this is the case. Let us assume we have two equations

$$\begin{aligned} t_1 &= t_2 t_3 \\ t_4 t_2 &= t_5 \end{aligned}$$

and we want to derive a property for t_1 . Properties of t_3 , t_4 and t_5 are known (colored green). As the properties of t_2 are not known as well, we have to use the second equation to proceed. If t_4 is invertible, we can solve to

$$t_2 = t_4^{-1} t_5.$$

Instead of substituting t_2 in $t_1 = t_2 t_3$, yielding

$$t_1 = t_4^{-1} t_5 t_3,$$

and then reason about properties of $t_4^{-1} t_5$, we first derive all properties of $t_4^{-1} t_5$, which are also properties of t_2 . In the final step, we derive all properties of $t_2 t_3$, obtaining the same properties for t_1 we would find by plugging one equation into the other.

Termination

The disadvantage of this approach is that it may not terminate either, and increasingly long properties are derived. This is not unexpected, since this approach aims

at replicating the process of substituting equations. The difference of the presented method is that the set of equations is finite, its size does not even change anymore after the matrix inversion step.³ Furthermore, for most iterative methods, no properties of products of more than three quantities are used. While it might be possible to use significantly longer properties to solve equations, most likely, they result in algorithms that use unnecessarily large expressions to compute certain quantities. This naturally leads to the solution of introducing a maximum length for properties, similar to a recursion limit, with a reasonable default value that can be changed by the user. This way, there is only a finite number of properties that can be derived.

Unfortunately, if we just refrained from adding properties of products of more than three quantities to \mathcal{P} , the derivation process would cease to work in certain cases (this will be explained in the following section). To avoid that, there are multiple options. In both cases, we initially add longer properties to \mathcal{P} as well. Then, one solution is to never use them to derive further properties, that is, we never construct $S(t)$ if t is a product of more than three quantities. Alternatively, those longer properties are removed from \mathcal{P} when e' is discarded. A third option would be to construct an additional set of temporary properties.

3.2.5. Example: Nonsymmetric CG

We demonstrate the method presented in this chapter by deriving some properties for nonsymmetric CG. In the interest of brevity, we only derive a small number of selected properties, in addition to limiting properties to a maximum length of three quantities. The knowledgebase \mathcal{K} is not shown here due to its size. Properties used in the following which are not self-explanatory are defined in Appendix A. The pre- and postcondition of nonsymmetric CG are shown below. For the sake of simplicity, we treat $\underline{I} - \underline{J}$ as one distinct matrix, as \underline{I} and \underline{J} do not appear separately.

$$\begin{aligned} P_{\text{pre}} : & \{ \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\ & \text{Output}[P] \wedge \text{Matrix}[P] \wedge \\ & \text{Output}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge \\ & \text{FirstColumnInput}[R] \wedge \text{Matrix}[R] \wedge \text{Orthogonal}[R] \wedge \\ & \text{FirstColumnInput}[\underline{R}] \wedge \text{Matrix}[\underline{R}] \wedge \text{Orthogonal}[\underline{R}] \wedge \\ & \text{Diagonal}[R[\underline{R}^T \underline{R}]] \wedge \text{Diagonal}[\underline{R}[\underline{R}^T R]] \wedge \\ & \text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \wedge \\ & \text{Output}[U] \wedge \text{Matrix}[U] \wedge \text{StrictlyUpperTriangular}[U] \wedge \\ & \text{Matrix}[\underline{I} - \underline{J}] \wedge \text{LowerTrapezoidal}[\underline{I} - \underline{J}] \} \end{aligned}$$

³This is why we refrain from adding e' to \mathcal{E} in “Application of Properties”, step 3.

3. Derivation of Algorithms for Iterative Methods: Foundations

$$\begin{aligned} P_{\text{post}} : \{ & \text{APD} = \mathbf{R} (\underline{\mathbf{I}} - \underline{\mathbf{J}}) \\ & \mathbf{P} (\mathbf{I} - \mathbf{U}) = \underline{\mathbf{R}} \\ & \mathbf{PD} = \mathbf{X} (\underline{\mathbf{I}} - \underline{\mathbf{J}}) \\ & \|\mathbf{Re}_r^T\| < \varepsilon \} \end{aligned}$$

Initialization

The first step consists of initializing \mathcal{P} and \mathcal{E} . The former contains all the properties of the precondition P_{pre} that describe operand types.

$$\begin{aligned} \mathcal{P} = \{ & \text{Matrix}[\mathbf{A}], \text{NonSingular}[\mathbf{A}], \\ & \text{Matrix}[\mathbf{P}], \\ & \text{Matrix}[\mathbf{D}], \text{Diagonal}[\mathbf{D}], \\ & \text{Matrix}[\mathbf{R}], \text{Orthogonal}[\mathbf{R}], \\ & \text{Matrix}[\underline{\mathbf{R}}], \text{Orthogonal}[\underline{\mathbf{R}}], \\ & \text{Diagonal}[\mathbf{R}^T \underline{\mathbf{R}}], \text{Diagonal}[\underline{\mathbf{R}}^T \mathbf{R}] \wedge \\ & \text{Matrix}[\mathbf{X}], \\ & \text{Matrix}[\mathbf{U}], \text{StrictlyUpperTriangular}[\mathbf{U}], \\ & \text{Matrix}[\underline{\mathbf{I}} - \underline{\mathbf{J}}], \text{LowerTrapezoidal}[\underline{\mathbf{I}} - \underline{\mathbf{J}}] \} \end{aligned}$$

\mathcal{E} contains the equations of the postcondition P_{post} .

$$\begin{aligned} \mathcal{E} = \{ & \text{APD} = \mathbf{R} (\underline{\mathbf{I}} - \underline{\mathbf{J}}), \\ & \mathbf{P} (\mathbf{I} - \mathbf{U}) = \underline{\mathbf{R}}, \\ & \mathbf{PD} = \mathbf{X} (\underline{\mathbf{I}} - \underline{\mathbf{J}}) \} \end{aligned}$$

Derivation of Properties

During this step, only a small number of new properties can be derived. The implication

$$\text{StrictlyUpperTriangular}[\mathbf{t}_1] \wedge \mathbf{I} + \mathbf{t}_1 \rightarrow \text{UpperTriangular}[\mathbf{I} + \mathbf{t}_1]$$

is used to infer that $(\mathbf{I} - \mathbf{U})$ is upper triangular. $\text{Orthogonal}[\mathbf{R}]$ and $\text{Orthogonal}[\underline{\mathbf{R}}]$ imply that $\mathbf{R}^T \mathbf{R}$ and $\underline{\mathbf{R}}^T \underline{\mathbf{R}}$ are diagonal. Thus, this step yields

$$\begin{aligned} \mathcal{P} = \mathcal{P} \cup \{ & \text{Diagonal}[\mathbf{R}^T \mathbf{R}], \\ & \text{Diagonal}[\underline{\mathbf{R}}^T \underline{\mathbf{R}}], \\ & \text{UpperTriangular}[\mathbf{I} - \mathbf{U}] \}. \end{aligned}$$

Matrix Inversion

From the set of properties \mathcal{P} it follows that A , D and $(I - U)$ are nonsingular. The equation $APD = R(\underline{I} - \underline{J})$ is inspected first. Two invertible objects, namely A and D , occur in it. Since A is the leftmost quantity of the product APD , it can be eliminated by multiplying its inverse from the left-hand side to both sides of the equation. This yields $PD = A^{-1}R(\underline{I} - \underline{J})$, which is added to \mathcal{E} . By multiplying D^{-1} from the right to this new equation and the original $APD = R(\underline{I} - \underline{J})$, two additional equations are obtained. By applying the same procedure to the two remaining equations, \mathcal{E} becomes the following set:

$$\begin{aligned} \mathcal{E} = \{ & APD = R(\underline{I} - \underline{J}), & AP &= R(\underline{I} - \underline{J})D^{-1}, \\ & PD = A^{-1}R(\underline{I} - \underline{J}), & P &= A^{-1}R(\underline{I} - \underline{J})D^{-1}, \\ & P(I - U) = \underline{R}, & P &= \underline{R}(I - U)^{-1}, \\ & PD = X(\underline{I} - \underline{J}), & P &= X(\underline{I} - \underline{J})D^{-1} \} \end{aligned}$$

At this point, it is not possible to find any new equations by multiplying inverted quantities, so we proceed to the next step.

Application of Properties

Initially, the only properties $P[t] \in \mathcal{P}$ with $t = t_1 \cdots t_i$, $i > 1$ are

$$\begin{aligned} & \text{Diagonal}[R^T R] & \text{Diagonal}[\underline{R}^T \underline{R}] \\ & \text{Diagonal}R[R^T \underline{R}] & \text{Diagonal}R[\underline{R}^T R]. \end{aligned}$$

From $\text{Diagonal}[R^T R]$, the set $S(R^T R) = \{(R^T, R)\}$ is obtained. For every $(t_L, t_R) \in S(R^T R)$, it now has to be checked if t_L or t_R appears in any equation contained in \mathcal{E} . Since (R^T, R) is the only element in $S(R^T R)$, the system just searches for $t_L = R^T$ and $t_R = R$. The following four equations, all containing R , are found:

$$\begin{aligned} APD &= R(\underline{I} - \underline{J}) & AP &= R(\underline{I} - \underline{J})D^{-1} \\ PD &= A^{-1}R(\underline{I} - \underline{J}) & P &= A^{-1}R(\underline{I} - \underline{J})D^{-1} \end{aligned}$$

Only in the first two equations, $t_L = R$ appears at the rightmost position in a product. Hence, $t_R = R^T$ is multiplied just to those two:

$$R^T APD = R^T R(\underline{I} - \underline{J}) \quad R^T AP = R^T R(\underline{I} - \underline{J})D^{-1}$$

The knowledgebase is now used to derive new properties. In this example, we just look at the equation on the right-hand side, and show the steps that the system would perform.

3. Derivation of Algorithms for Iterative Methods: Foundations

1. $R^T R (\underline{I} - \underline{J})$ is a product of a diagonal and a lower trapezoidal matrix, so it is lower trapezoidal.
2. D is diagonal, so D^{-1} is diagonal as well.
3. $R^T R (\underline{I} - \underline{J}) D^{-1}$ is a product of a lower trapezoidal matrix ($R^T R (\underline{I} - \underline{J})$) and a diagonal matrix (D^{-1}), so it is lower trapezoidal too.
4. The right-hand side of the equation is lower trapezoidal, so the left-hand side of the equation, $R^T A P$, is lower trapezoidal as well.

With every step, the new properties are added to \mathcal{P} . Thus, the set becomes:

$$\begin{aligned} \mathcal{P} := \mathcal{P} \cup \{ & \text{LowerTrapezoidal}[R^T R (\underline{I} - \underline{J})], \\ & \text{Diagonal}[D^{-1}], \\ & \text{LowerTrapezoidal}[R^T R (\underline{I} - \underline{J}) D^{-1}], \\ & \text{LowerTrapezoidal}[R^T A P] \} \end{aligned}$$

Now, it becomes apparent why it is not possible to simply set a limit on the length of properties and never add any longer properties to \mathcal{P} . $R^T R (\underline{I} - \underline{J}) D^{-1}$ is a product of four quantities. To infer that $R^T A P$ is lower trapezoidal, the property

$$\text{LowerTrapezoidal}[R^T R (\underline{I} - \underline{J}) D^{-1}]$$

has to be derived. If this property is never added to \mathcal{P} , it is not possible to derive that $R^T A P$ is lower trapezoidal. Hence, it must be possible to derive properties of any length, even though they are only needed temporarily.

Using the newly added properties, immediately some more are found because of the implications

$$\begin{aligned} \text{LowerTrapezoidal}[t] &\rightarrow \text{UpperTrapezoidal}[t^T] \\ (t_1 t_2)^T &\rightarrow t_2^T t_1^T. \end{aligned}$$

The following properties are added to the set:

$$\begin{aligned} \mathcal{P} := \mathcal{P} \cup \{ & \text{UpperTrapezoidal}[(\underline{I} - \underline{J})^T R^T R], \\ & \text{Diagonal}[D^{-T}], \\ & \text{UpperTrapezoidal}[D^{-T} (\underline{I} - \underline{J}) R^T R], \\ & \text{UpperTrapezoidal}[P^T A R] \} \end{aligned}$$

At this point, we cut this derivation short. In practice, many more properties would be derived. Most of them may not be of any use for the derivation of algorithms, but a few are crucial. For nonsymmetric CG, such an important property is that $P^T A P$ is lower triangular. In [5], this property is derived manually. Let us shortly illustrate how it is derived.

1. Because of `DiagonalR[RTR]`, R^T is multiplied from the left to $P = \underline{R} (I - U)^{-1}$, yielding

$$R^T P = R^T \underline{R} (I - U)^{-1}.$$

$(I - U)^{-1}$ is upper triangular and $R^T \underline{R}$ is rectangular diagonal, so $R^T P$ is upper triangular and rectangular. It follows that its transpose $P^T R$ is lower triangular and rectangular.

2. Because of the property `LowerTriangularR[PTR]`, P^T is multiplied from the left to $APD = R (\underline{I} - \underline{J})$. Based on the resulting equation

$$P^T APD = P^T R (\underline{I} - \underline{J}),$$

we derive that $P^T R (\underline{I} - \underline{J})$ is square.

3. For the same reason, P^T is multiplied from the left to $AP = R (\underline{I} - \underline{J}) D^{-1}$, resulting in

$$P^T AP = P^T R (\underline{I} - \underline{J}) D^{-1}.$$

$P^T R$ is lower triangular and rectangular and $\underline{I} - \underline{J}$ is lower trapezoidal. $P^T R (\underline{I} - \underline{J})$ is square, so it is lower triangular. Since D^{-1} is diagonal, $P^T AP$ is lower triangular as well.

3.3. Initial Partitionings

The first step towards the derivation of algorithms in the FLAME methodology is to partition the operands. The matrix representation of iterative methods gives rise to a new type of operands, namely those which are initially partially known and partially unknown. Thus, we have to address the question of how to partition them. To some extent, it is easy to answer. Standard CG algorithms always compute full vectors r_i and x_i [1, 11, 13], so refraining from partitioning R and X horizontally, that is, into a top and a bottom part, is a very natural choice. For partitioning vertically, there are multiple possibilities. Their advantages and disadvantages will be discussed in this section.

While not strictly necessary to derive algorithms, describing the partitioned operands in terms of functions proved to be useful for the systematic derivation [7]. Hence, if and how different partitionings permit to match functions is an important criterion for their evaluation.

We begin this section with a discussion of what such a function should look like. Then, we investigate different possible partitionings, starting with partitionings that are also used for direct methods and continuing with some that are tailored to properties of algorithms for iterative methods.

Functions

On the highest level, the input of a CG algorithm are vectors b and x_0 , as well as the matrix A , so a function would have the form $x_i := \text{CG}(A, b, x_0)$. As we are using the recurrence relation to derive algorithms, which does not involve b , but a number of other quantities, for example the residual r , a function like this is not helpful.

From direct methods, we remember that the function is uniquely⁴ defined by the precondition. Every quantity that is initially known has the property **Input**[X], initially unknown quantities have the property **Output**[X]. Take the LU factorization as an example. The governing equation is $LU = A$. A is input, L and U are output. This naturally leads to a function $\{L, U\} = \Psi(A)$. It is important to note that this function is defined prior to any derivation steps, solely based on the abstract description of the operation. It then happens, due to the recursive nature of the operations, *and a partitioning that reveals it*, that this function matches expressions obtained by partitioning the postcondition. In some cases, the expressions have to be rewritten first. The remaining expressions can be decomposed into basic buildings blocks.

Applying this scheme to CG, the function $\{R, U, P, D, X\} := \text{CG}(A, R_{e0}, X_{e0})$ is obtained. For simplicity, we omit I and J as input, since they are constant.

Standard 2 x 2 Partitioning

Naturally, the first choice for a partitioning is the one that is usually used for direct methods. Partitioning R into $(R_L \mid R_R)$ implies the following partitioning for equations (3.1) and (3.2). Because of its similarity to the first equation, we will usually omit equation (3.3) in this section.

$$\begin{aligned} A (P_L \mid P_R) \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & D_{BR} \end{array} \right) &= (R_L \mid R_R) \left(\begin{array}{c|c} I-J & 0 \\ \hline -H & I-J \end{array} \right) \\ (P_L \mid P_R) \left(\begin{array}{c|c} I-U_{TL} & -U_{TR} \\ \hline 0 & I-U_{BR} \end{array} \right) &= (R_L \mid \underline{R}_R) \end{aligned}$$

Here, H is a matrix with one more row than columns that is one in the top right corner and zero everywhere else. Flattening the expressions, we obtain

$$(AP_L D_{TL} = R_L (I-J) - R_R H \mid AP_R D_{BR} = R_R (I-J)) \quad (3.12)$$

$$(P_L (I-U_{TL}) = R_L \mid -P_L U_{TR} + P_R (I-U_{BR}) = \underline{R}_R) \quad (3.13)$$

In this form, the expressions are not matched by the pattern of the CG function. While both the right equation of (3.12) and the left one of (3.13) have the correct shape, one contains R_L and P_L , while R_R and P_R appear in the other. In the original equations, those quantities are the same. If it is possible to match the function at all, then either

⁴Except for the ordering of input and output.

all parts on the left or all parts on the right match (or both). Rewriting the equation on the left in (3.12) as

$$AP_L D_{TL} = (R_L \mid R_R) \left(\begin{array}{c} I-J \\ -H \end{array} \right),$$

we recognize similarities to the corresponding equation of the recurrence relation (3.1). However,

$$\left(\begin{array}{c} I-J \\ -H \end{array} \right)$$

does not have one more row than columns. While we know that only the first row of H has a nonzero entry, formally, this equation does not match the pattern.

3 x 3 Partitioning

The problem above can be solved by applying a 1×3 partitioning to R where the middle part is a single column. This has the effect that the first row of H becomes a separate block. As CG proceeds by one column per iteration, exposing a single column appears to be a suitable choice. The partitioned operands and the resulting expressions are shown below.

$$\begin{aligned} A (P_L \mid p_M \mid P_R) \left(\begin{array}{c|c|c} D_{TL} & 0 & 0 \\ \hline 0 & \delta_{MM} & 0 \\ \hline 0 & 0 & D_{BR} \end{array} \right) &= (R_L \mid r_M \mid R_R) \left(\begin{array}{c|c|c} I-J & 0 & 0 \\ \hline -e_r^T & 1 & 0 \\ \hline 0 & -e_0 & I-J \end{array} \right) \\ (P_L \mid p_M \mid P_R) \left(\begin{array}{c|c|c} I-U_{TL} & -u_{TM} & -u_{TR} \\ \hline 0 & 1 & -u_{MR} \\ \hline 0 & 0 & I-U_{BR} \end{array} \right) &= (R_L \mid r_M \mid \underline{R}_R) \end{aligned}$$

$$\begin{aligned} (AP_L D_{TL} = R_L (I-J) - r_M e_r^T \mid Ap_M \delta_{MM} = r_M - R_R e_0 \mid AP_R D_{BR} = R_R (I-J)) \\ (P_L (I-U_{TL}) = R_L \mid -P_L u_{TM} + p_M = r_M \mid -P_L u_{TR} - p_M u_{MR} + P_R (I-U_{BR}) = \underline{R}_R) \end{aligned}$$

Now, a similar rewriting as in the previous section yields

$$AP_L D_{TL} = (R_L \mid r_M) \left(\begin{array}{c} I-J \\ -e_r^T \end{array} \right).$$

Here, the constant matrix has one more row than columns, and $(R_L \mid r_M)$ has one more column than R_L in $P_L (I-U_{TL}) = R_L$. Hence, those equations, together with the additional one for X , are matched by the pattern of the CG function. The resulting

3. Derivation of Algorithms for Iterative Methods: Foundations

assignment is

$$\{ (\mathbf{R}_L \mid \mathbf{r}_M), \mathbf{U}_{TL}, \mathbf{P}_L, \mathbf{D}_{TL}, (\mathbf{X}_L \mid \mathbf{x}_M) \} := \text{CG} \left(\mathbf{A}, (\mathbf{R}_L \mid \mathbf{r}_M) \left(\frac{\mathbf{e}_0}{0} \right), (\mathbf{X}_L \mid \mathbf{x}_M) \left(\frac{\mathbf{e}_0}{0} \right) \right).$$

While the equations in the middle are not matched by the function, they can be solved to \mathbf{u}_{TM} , \mathbf{p}_M , δ_{MM} and $\mathbf{R}_R \mathbf{e}_0$, leading to computable assignments. Unfortunately, the function does not match the equations on the right-hand side. For direct methods, in a comparable situation, an equation like

$$-\mathbf{P}_L \mathbf{U}_{TR} - \mathbf{p}_M \mathbf{u}_{MR} + \mathbf{P}_R (\mathbf{I} - \mathbf{U}_{BR}) = \underline{\mathbf{R}}_R$$

would be rewritten as

$$\mathbf{P}_R (\mathbf{I} - \mathbf{U}_{BR}) = \underline{\mathbf{R}}_R + \mathbf{P}_L \mathbf{U}_{TR} + \mathbf{p}_M \mathbf{u}_{MR},$$

such that $\underline{\mathbf{R}}_R$ is updated. Here, this is not possible. While the first column of $\underline{\mathbf{R}}_R$ can be considered known at this point, neither the first column of $\mathbf{P}_L \mathbf{U}_{TR}$ nor $\mathbf{p}_M \mathbf{u}_{MR}$ is known, because neither \mathbf{U}_{TR} nor \mathbf{u}_{MR} are known. Apart from that, $\underline{\mathbf{R}}_R$ can not be updated, as this would have influences on the other two equations. Finally, if quantities are updated, they are usually updated in their entirety, before they are used as input for a function. Here, \mathbf{U}_{TR} and \mathbf{u}_{MR} are not known, and we would expect them to be the output of said function, leading to circular data dependencies.

Splitting off the First Column

Since the first column of \mathbf{R} and \mathbf{X} , respectively, plays a special role, an entirely different approach could be to apply a partitioning that splits off this first column. \mathbf{R} is partitioned into $(\mathbf{r}_0 \mid \mathbf{R}')$, and for the actual derivation of algorithms, \mathbf{R}' is used. The advantage is that there is a clear distinction between input and output, and it would be possible to write CG as $\{ \mathbf{R}', \dots, \mathbf{X}' \} := \text{CG}(\mathbf{A}, \mathbf{r}_0, \mathbf{x}_0)$. As usual, partitioning \mathbf{R} and \mathbf{X} like that also implies a matching partitioning for the remaining operands:

$$\begin{aligned} \mathbf{A} (\mathbf{p}_0 \mid \mathbf{P}') \left(\frac{\delta_0 \mid 0}{0 \mid \mathbf{D}'} \right) &= (\mathbf{r}_0 \mid \mathbf{R}') \left(\frac{1 \mid 0}{-\mathbf{e}_0 \mid \mathbf{I} - \mathbf{I}} \right) \\ (\mathbf{p}_0 \mid \mathbf{P}') \left(\frac{1 \mid -\mathbf{u}'}{0 \mid \mathbf{I} - \mathbf{U}'} \right) &= (\mathbf{r}_0 \mid \underline{\mathbf{R}}') \end{aligned}$$

Flattening those expression, we immediately obtain a value for \mathbf{p}_0 :

$$\begin{aligned} (\mathbf{A} \mathbf{p}_0 \delta_0 = \mathbf{r}_0 - \mathbf{R}' \mathbf{e}_0 \mid \mathbf{A} \mathbf{P}' \mathbf{D}' = \mathbf{R}' (\mathbf{I} - \mathbf{I})) \\ (\mathbf{p}_0 = \mathbf{r}_0 \mid -\mathbf{p}_0 \mathbf{u}' + \mathbf{P}' (\mathbf{I} - \mathbf{U}') = \underline{\mathbf{R}}') \end{aligned}$$

Using the orthogonality of R , it is also possible to find an assignment for δ_0 . Thus, this partitioning allows to compute all quantities of the first iteration and declare them as known. Unfortunately, splitting off r_0 forces us to also split off the first row of U . At this point, it is not possible to compute it in its entirety. Furthermore, because of $-p_0 u'$, the equations on the right-hand side do not have the same shape as the original description of CG. While it is of course still possible to compute u' , it is only possible entry by entry, adding an additional assignment to any update we can derive. Consequently, it will not be possible to derive the updates for CG usually found in textbooks.

Divide and Conquer

So far, no partitioning enabled us to describe one CG operation as multiple, smaller CG operations, if necessary, with updated quantities as input, similar to how the PME of the triangular system in Section 2.1 contains the function Φ two times. This is, however, possible, if we define a generalized version of CG. The disadvantage is that it requires a deeper understanding of the algorithm, in addition to some knowledge that is initially not available when deriving algorithms solely based on their matrix representation.

To derive such a representation, we partition R into $(R_0 \mid r_1 \mid r_2 \mid R_3)$. Since the resulting partitioned postcondition is very large, it is not shown here. The generalized version of CG requires some parts of P as an additional argument. Initially, it is only the first column P , here denoted with P_{e_0} : $\{R, U, P, D, X\} := \text{CG}(A, R_{e_0}, X_{e_0}, P_{e_0})$. Since it is equal to the first column of R , it can be considered known. For now, we will assume that A is nonsymmetric.

Our goal is now to write this CG operation as two separate ones, one covering $(R_0 \mid r_1)$, and one for $(r_2 \mid R_3)$. Clearly, the first one is

$$\{ (R_0 \mid r_1 \mid r_2), \dots, (P_0 \mid p_1), \dots, (X_0 \mid x_1 \mid x_2) \} := \text{CG}(A, R_0 e_0, X_0 e_0, P_0 e_0),$$

omitting some of the output in the interest of legibility. Now, to find the correct arguments for the second one, we need to know how p_2 is computed:

$$p_2 = r_2 + P_0 u_{02} + p_1 v_{12}$$

u_{02} and v_{12} are not part of the output of the function above, but they can in turn be computed using known quantities only:

$$u_{02} = \left(-P_0^T A P_0 \right)^{-1} \cdot P_0^T A r_2$$

$$v_{12} = -\frac{p_1^T A r_2 + p_1^T A P_0 u_{02}}{p_1^T A p_1}$$

3. Derivation of Algorithms for Iterative Methods: Foundations

To eliminate them entirely, we can also write

$$p_2 = r_2 - (P_0 \mid p_1) \left(\frac{P_0^T A P_0 \mid 0}{p_1^T A P_0 \mid p_1^T A p_1} \right)^{-1} \left(\frac{P_0^T A r_2}{p_1^T A r_2} \right).$$

Now, we can write the second part as

$$\{R_3, \dots, P_3, \dots, X_3\} := \text{CG}(A, r_2, x_2, (P_0 \mid p_1 \mid p_2)).$$

To ensure that this function computes the correct sequence of search directions, it also has to use the ones computed by the first function, which are $(P_0 \mid p_1)$. This is the reason why not just p_2 , but $(P_0 \mid p_1 \mid p_2)$ is input.

This situation is slightly different if A is symmetric. In that case, U is upper diagonal and $P^T A P$ is diagonal, so p_2 is computed as

$$p_2 = r_2 - p_2 \frac{p_1^T A r_2}{p_1^T A p_1}.$$

Now, each search direction is computed using only the last one, so $(P_0 \mid p_1)$ is not needed as input. Thus, the second function simplifies to

$$\{R_3, \dots, P_3, \dots, X_3\} := \text{CG}(A, r_2, x_2, p_2).$$

As mentioned before, the disadvantage is that we already need to know how some quantities are computed to derive this representation, while it is actually our goal to find those updates.

Splitting off the Last Column

Among those presented in this section, the 3×3 partitioning that exposes a single column is the only one that resulted in expressions that were naturally matched by the CG function. The problem of this partitioning is that there is no obvious way how to deal with the right-hand side parts.

To find one that better suits iterative methods, it is helpful to again inspect the differences to direct methods. After all, the 3×3 partitioning came to existence as a modification of the standard partitioning used for direct methods. As mentioned before, with direct methods, the sizes of all operands are initially known. Thus, at any point during the computation, there are (potentially empty) parts of operands that are already computed, and (potentially empty) parts that are not computed yet. The loop invariant describes those parts that are already computed at the beginning of the loop body. Conversely, one can think of those parts of the PME that are not part of the loop invariant as those parts that are not computed yet.

To see that this is consistent, recall that the reason that the full set of nodes can never be a feasible loop invariant is that it would imply that the solution is already

computed before the loop is entered. Clearly, this is equivalent to saying that there are no parts left that are not computed yet, even before the loop is entered.

In case of iterative methods, there is little use in talking about parts of operands that are not yet computed beyond the current iteration, as each iteration might be the last. We can conclude that it makes little sense to use a partitioning where the right-hand side is more than a single column.

The solution is to use a partitioning that is a hybrid of the standard 2×2 partitioning and the 3×3 partitioning that exposes a single column: A partitioning that splits off the last column: P is partitioned into $(P_L \mid p_R)$. Because of the additional column of R , it is partitioned into $(R_L \mid r_R \mid r_+)$. Thus, for CG, we obtain

$$\begin{aligned} A (P_L \mid p_R) \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & \delta_{BR} \end{array} \right) &= (R_L \mid r_R \mid r_+) \left(\begin{array}{c|c} I-J & 0 \\ \hline -e_r^T & 1 \\ \hline 0 & -1 \end{array} \right) \\ (P_L \mid p_R) \left(\begin{array}{c|c} I-U_{TL} & -u_{TR} \\ \hline 0 & 1 \end{array} \right) &= (R_L \mid r_R) \\ (P_L \mid p_R) \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & \delta_{BR} \end{array} \right) &= (X_L \mid x_R \mid x_+) \left(\begin{array}{c|c} I-J & 0 \\ \hline -e_r^T & 1 \\ \hline 0 & -1 \end{array} \right). \end{aligned}$$

Flattening those expressions yields

$$\begin{aligned} (AP_L D_{TL} = R_L(I-J) - r_R e_r^T \mid Ap_R \delta_{BR} = r_R - r_+) \\ (P_L(I-U_{TL}) = R_L \mid -P_L u_{TR} + p_R = r_R) \\ (P_L D_{TL} = X_L(I-J) - x_R e_r^T \mid p_R \delta_{BR} = x_R - x_+). \end{aligned}$$

The left-hand side parts are now matched by

$$\begin{aligned} \{ (R_L \mid r_R), U_{TL}, P_L, D_{TL}, (X_L \mid x_R) \} := \\ \text{CG} \left(A, (R_L \mid r_R) \left(\frac{e_0}{0} \right), (X_L \mid x_R) \left(\frac{e_0}{0} \right) \right). \end{aligned}$$

Using some of the properties derived with the approach presented in Section 3.2, the equations on the right can be solved to all remaining unknowns. Thus, we obtain a PME with assignments for every unknown quantity. The systematic derivation of loop-based algorithms, using this partitioning, is presented in the next chapter.

Note that this PME can be interpreted as an “inductive PME”: Assuming it is possible to compute an arbitrary number of previous iterations, it is possible to compute one additional iteration. The previous iterations are represented by the CG function, and the additional iteration is computed using the remaining, explicit assignments of the PME. The base case is obtained by assuming all left and top left parts to be empty.

4. Derivation of Algorithms for Iterative Methods

After having laid the foundations in the previous chapter, the actual approach for deriving algorithms for iterative methods is presented in this chapter. The approach itself can be found in Section 4.1, followed by two examples in Section 4.2 and 4.3. In the final Section 4.4, scope and limitations of the approach are discussed.

4.1. Derivation of Algorithms

The derivation of algorithms for iterative methods mainly follows the same three basic steps as for direct methods. First, one or more PME's are generated. In the second step, dependency graphs are constructed, which are then used to select loop invariants. In the third and final step, from each loop invariant, an algorithm is constructed. In this section, we will present this process for iterative methods. There is, however, a new fourth step. In this step, some postprocessing is applied to the derived algorithms to generate a number of variants that may behave differently in floating point arithmetic or vary in their performance. If applicable, we follow the same structure as [6].

4.1.1. PME Generation

The first stage towards the generation of algorithms is to find PME's. The necessary steps are explained in the following. There are a number of differences compared to direct methods: It is necessary to derive properties of matrices, which then enable the system to solve equations. Additionally, different partitionings are needed to deal with new types of operands.

Derivation of Properties

The approach presented in Section 3.2 is used to derive properties.

Initial Partitioning

Operands of the postcondition are partitioned depending on their shape and properties. For objects that are completely known or unknown, the applicable partitionings

4. Derivation of Algorithms for Iterative Methods

are similar to the ones used for direct methods. The difference lies in the sizes of the resulting objects. Only the top, left, and top left parts are matrices, the remaining ones are either vectors or scalars:

$$\begin{array}{ll} B \rightarrow B & B \rightarrow (B_L \mid b_R) \\ B \rightarrow \left(\frac{B_T}{b_B} \right) & B \rightarrow \left(\frac{B_{TL} \mid b_{TR}}{b_{BL} \mid \beta_{BR}} \right) \end{array}$$

Just as with direct methods, triangular or symmetric matrices are either not partitioned at all, or the 2×2 partitioning is used. In case of the latter, the top left part is required to be square, such that it inherits the property of the matrix.

Separate partitionings are necessary to deal with matrices where initially, only the first column is known. If the last column of those matrices is omitted, the usual 1×2 partitioning is applied:

$$\underline{B} \rightarrow (B_L \mid b_R)$$

In case of the complete matrix, an additional column is obtained:

$$B \rightarrow (B_L \mid b_R \mid b_+)$$

The constant matrices \underline{J} and \underline{I} pose a special case. They have one more row than columns, so \underline{J} is not lower diagonal and \underline{I} is not an identity matrix. To derive algorithms, it is important to utilize their specific structure, so we will provide the partitionings explicitly:

$$\underline{J} \rightarrow \left(\frac{J \mid 0}{e_r^T \mid 0} \right) \quad \underline{I} \rightarrow \left(\frac{I \mid 0}{0 \mid 1} \right)$$

The same partitioning that is applied to the operands in the postcondition is also applied to the operands in the set of properties. Then, properties of expressions of partitioned operands are derived, similar to how partitioned operands inherit properties. If for example B is partitioned into $(B_L \mid b_R)$, and $\text{Diagonal} [B^T B]$ is contained in the set of properties, the system obtains

$$\text{Diagonal} \left[\left(\frac{B_L^T}{b_R^T} \right) (B_L \mid b_R) \right] = \text{Diagonal} \left[\left(\frac{B_L^T B_L \mid B_L^T b_R}{b_R^T B_L \mid b_R^T b_R} \right) \right].$$

Thus, it is possible to derive that $B_L^T B_L$ is diagonal as well, and $B_L^T b_R$ and $b_R^T B_L$ are zero.

Finding the PME

The first part of this step consists of performing symbolic arithmetic and distributing equalities across the partitionings. Consider the Krylov sequence as an example:

$$K := KS(A, Ke_0) \equiv \begin{cases} P_{\text{pre}} : \{ \text{Input}(A) \wedge \text{Matrix}(A) \\ \quad \text{Matrix}[J] \wedge \text{LowerDiagonalR}[J] \wedge \\ \quad \text{FirstColumnInput}(K) \} \\ P_{\text{post}} : \{ AK = KJ \\ \quad \text{size}(K) = n \times m \} \end{cases}$$

The partitioned postcondition looks as follows:

$$A \left(\begin{array}{c|c} K_L & k_R \end{array} \right) = \left(\begin{array}{c|c|c} K_L & k_R & k_+ \end{array} \right) \left(\begin{array}{c|c} J & 0 \\ \hline e_r^T & 0 \\ \hline 0 & 1 \end{array} \right)$$

It can be rewritten as the following expression:

$$\left(AK_L = K_L J + k_R e_r^T \mid Ak_R = k_+ \right)$$

In the second part, the goal is to find a representation of this expression where the value of each unknown quantity is determined by an assignment, using known operations. The quantities on the right-hand side of that assignment either have to be known, or their value is determined by another assignment. Intuitively, one could say the goal is to make this expression computable. This is done by matching patterns of known functions and operations.

Rewriting the equation on the left-hand side as

$$AK_L = \left(\begin{array}{c|c} K_L & k_R \end{array} \right) \left(\begin{array}{c} J \\ \hline e_r^T \end{array} \right),$$

it is easy to see that it describes the computation of a Krylov sequence, so it is possible to use the function from the description of the operation:

$$\left(\begin{array}{c|c} K_L & k_R \end{array} \right) := KS \left(A, \left(\begin{array}{c|c} K_L & k_R \end{array} \right) \left(\begin{array}{c} e_0 \\ \hline 0 \end{array} \right) \right)$$

Now, K_L and k_R can be considered known, so all that remains is to find an assignment for k_+ . In this example, the equation on the right hand side is already solved to k_+ , so

$$k_+ := Ak_R$$

4. Derivation of Algorithms for Iterative Methods

is immediately obtained. The PME then is

$$\left((K_L \mid k_R) := KS \left(A, (K_L \mid k_R) \left(\frac{e_0}{0} \right) \right) \mid k_+ := Ak_R \right).$$

In general, finding assignments is not that straightforward. Usually, it is necessary to apply known properties. Similar to the “Application of Properties” step in Section 3.2.3, parts of expressions with known properties are multiplied to both sides of equations, with the intention to recreate said expressions. Unfortunately, it is not easy to determine beforehand if the application of a property allows to solve an equation. Thus, in the manner of an exhaustive search, all matching properties are applied. If by this means, multiple assignments for the same expression are found, separate PME are derived for each variant.

Remark on Recursive Algorithms

The PME for direct methods immediately lead to recursive algorithms [2, 6]. This is also true for iterative methods. The difference is that in case of direct methods, those algorithms are usually divide and conquer algorithms. For iterative methods, the PME naturally leads to a “head recursive” implementation, that is, the recursive function call is the first operation in the function body.

4.1.2. Loop Invariant Identification

For the second step, the identification of loop invariants, constraints for the feasibility of loop invariants are introduced that differ from the ones used for direct methods. To understand why they are introduced, it is helpful to get an intuition for what those loop invariants express, and how they differ from the ones for direct methods.

Loop Invariants for Iterative Methods For iterative methods, there is usually a (partial) ordering in which parts of different operands can be computed. Let us assume there are three matrices B, C and D. During the execution of the algorithm, the following sequence is computed, where b_k , c_k and d_k are columns of B, C and D, respectively:

$$\dots b_i \ c_i \ d_i \ b_{i+1} \ c_{i+1} \ d_{i+1} \dots$$

A variant derived from one loop invariant may compute $c_i \ d_i \ b_{i+1}$ in one iteration. In contrast, one derived from another loop invariant may compute $d_i \ b_{i+1} \ c_{i+1}$. Since no quantities are overwritten, a loop invariant only expresses at which point in the sequence above an iteration starts.

In contrast, different loop invariants for direct methods may result in algorithms that compute the solution by row or by column.

Graph of Dependencies

The construction of the dependency graph is even simpler compared to direct methods. The assignments of the PME are not decomposed into basic building blocks. Each assignment is represented by one node in the dependency graph. The dependencies are established as usual.

Subset Selection

Just as with direct methods, subsets of nodes of the dependency graph are selected as candidates for loop invariants. Again, for every node that is contained in a subset, all preceding nodes have to be in that set, too. To assess the feasibility of loop invariants, however, slightly modified constraints have to be imposed:

1. There must exist a basic initialization of the operands, that is, an initial partitioning, followed by some preprocessing operations, that renders the predicate P_{inv} true:

$\{P_{\text{pre}}\}$
Partition
 Preprocessing
 $\{P_{\text{inv}}\}$

2. P_{inv} and the negation of the loop guard, G , must imply the postcondition, P_{post} :

$$P_{\text{inv}} \wedge \neg G \Rightarrow P_{\text{post}}$$

We begin with discussing the first condition. How the operands are partitioned was already established in the Section “Initial Partitioning”, but the initial sizes of the partitioned operands were not specified. The iterative methods covered in this thesis proceed through those matrices where initially, only the first column is known, from the left to the right. As a consequence, the left parts of those matrices are initially empty. The initial block sizes for all partitionings are given in Table 4.1.

Recall that for direct methods, the full set of nodes of the dependency graph can never be a feasible loop invariant. Formally, the reason is that there is no initial partitioning that renders this loop invariant true. Alternatively, one can say that it implies that the complete solution is already computed even before the loop is entered. For iterative methods, this is different. Due to the initial partitioning, which does not expose blocks that represent the remaining iterations, there is no subset implying that the entire solution is already computed. Thus, the full set is a feasible loop invariant.

However, further loop invariant candidates for iterative methods are not rendered true by the initial partitioning for a different reason. This is the case for some subsets that contain more than the first node. The reason is that the additional nodes may

4. Derivation of Algorithms for Iterative Methods

Initial Partitioning	Dimensions
$B \rightarrow (B_L \mid b_R)$	B_L is $n \times 0$
$B \rightarrow \left(\frac{B_T}{b_R} \right)$	B_T is $0 \times n$
$B \rightarrow \left(\frac{B_{TL} \mid b_{TR}}{b_{BL} \mid \beta_{BR}} \right)$	B_{TL} is 0×0
$\underline{B} \rightarrow (B_L \mid b_R)$	B_L is $n \times 0$
$B \rightarrow (B_L \mid b_R \mid b_+)$	B_L is $n \times 0$

Table 4.1.: Initial sizes of partitioned operands.

compute quantities that are not empty in the initial partitioning. Consider the Krylov sequence as an example: k_R and k_+ remain $n \times 1$ vectors in the initial partitioning. Thus, the assignment $k_+ := Ak_R$ computes a non-empty quantity, even if the initial partitioning is applied. However, at the beginning of the operation, only k_R is known, which is the first column of K in the initial partitioning, while k_+ is the second column, which is unknown.

Fortunately, the equations from those additional nodes can be rendered true by computing those quantities with some preprocessing operations. Naturally, the necessary preprocessing operations are obtained by applying the initial partitioning to the operations of those nodes. In this example, the preprocessing operation is $k_+ := Ak_R$. Those preprocessing operations can consist of the same type of operations as the actual update operations.

To demonstrate how the first constraint is checked, we return to the example of the Krylov sequence. We start with the following loop invariant candidate:

$$\left((K_L \mid k_R) := KS \left(A, (K_L \mid k_R) \left(\frac{e_0}{0} \right) \right) \mid \neq \right)$$

If K_L has size $n \times 0$, the expression on the left-hand side becomes $k_R := KS(A, k_R)$. Initially, the first column of K , which is now k_R , is known. Thus, since both sides of this assignment are known, this expression is considered to be true. This implies that the loop invariant satisfies the first constraint.

The second candidate for a loop invariant is

$$\left((K_L \mid k_R) := KS \left(A, (K_L \mid k_R) \left(\frac{e_0}{0} \right) \right) \mid k_+ := Ak_R \right).$$

Note that in the interest of simplicity, we usually just write $K_L e_0$ instead of

$$(K_L \mid k_R) \left(\frac{e_0}{0} \right),$$

if it is sufficient. The following expression is obtained if the initial partitioning is applied, where K_L is empty:

$$(k_R := KS(A, k_R) \mid k_+ := Ak_R)$$

As k_+ is not known, the initial partitioning alone does not render this expression true. This, however, can be solved with a preprocessing operation. Here, the operation is $k_+ := Ak_R$. It follows that this loop invariant satisfies the first constraint as well.

To check if the second condition for the feasibility of loop invariants is satisfied, we first need to determine the loop guard. As discussed in Section 3.1, comparing the size of a growing block of a matrix to the size of the entire matrix does not work, as the matrix grows as well. For this reason, the additional predicates were added to the postcondition (by hand), which are now easily translated into loop guards (automatically). How this is done is shown in Table 4.2.

Predicate	Loop guard	
	FirstColumnInput[B]	Output[B]
$\ Be_r^T\ < \varepsilon$	$\ b_R\ \geq \varepsilon$	$\ B_L e_r^T\ \geq \varepsilon$
$\text{size}(B) = n \times k$	$\text{size}((B_L \mid b_R)) < n \times k$	$\text{size}(B_L) < n \times k$
$\ Be_r^T - Be_{r-1}^T\ < \varepsilon$	$\ b_R - B_L e_r^T\ \geq \varepsilon$	$\ B_L e_r^T - B_L e_{r-1}^T\ \geq \varepsilon$

Table 4.2.: Look-up table for determining loop guards for iterative methods. The row is selected according to the additional predicate in the postcondition. The column is selected depending on the property of the operand that appears in the position of B in that predicate. Example: The predicate is $\|Be_r^T\| < \varepsilon$. The precondition contains the property FirstColumnInput[R]. Thus, the loop guard is $\|r_R\| \geq \varepsilon$. To allow for other loop guards, this table has to be extended manually.

Note that even though Be_r^T in the predicate $\|Be_r^T\| < \varepsilon$ refers to the last column of B, b_R in the loop guard $\|b_R\| \geq \varepsilon$ and $B_L e_r^T$ in $\|B_L e_r^T\| \geq \varepsilon$, respectively, are the second to last columns. Similarly, the other loop guards omit the last column as well. To understand why this is necessary, we have to look at the loop invariant candidates again. Just as with direct methods, the empty set can never be a valid loop invariant because it corresponds to the empty predicate. Thus, all remaining candidates contain at least the first node of the dependency graph. For iterative methods, this first node

4. Derivation of Algorithms for Iterative Methods

always represents the operation itself, that is, it contains the original function. Which parts of the operands are output of that function depends on the properties of the operands:

- If initially, the first column of B is known ($\text{FirstColumnInput}[B]$), B_L and b_R are output of the function.
- If B is initially unknown ($\text{Output}[B]$), just B_L is output.

Consider nonsymmetric CG as an example: The operation of the first node is

$$\{ (R_L \mid r_R), U_{TL}, P_L, D_{TL}, (X_L \mid x_R) \} := \text{CG}(A, R_L e_0, X_L e_0).$$

The first columns of R and X, here denoted by $R_L e_0$ and $X_L e_0$, are initially known and the function computes R_L , r_R , X_L and x_R . In contrast, the property of P is $\text{Output}[P]$, and only P_L is computed. Intuitively, the reason is that the function always computes the same number of columns of all operands. If initially, one column of an operand is already known, in the end, one additional column is obtained.

Since only the first node is guaranteed to be part of the loop invariant, only those parts computed in the first node are guaranteed to be known at the beginning and at the end of the loop. Thus, for $\text{FirstColumnInput}[B]$, checking the norm of b_R in the loop guard is always possible, but using b_+ is not.

While there are variants which compute b_+ (or b_R in case of $\text{Output}[B]$), and those would allow different loop guards, we refrain from using those to keep the derivation simple. This, however, means that if b_+ (or b_R) are computed, they will not be regarded as part of the solution. That is possible because B does not have a fixed size. Unfortunately, this has the effect that some algorithms compute results that are subsequently discarded. If, however, the loop body of those algorithms has a reduced computational complexity, this is an acceptable tradeoff.

We demonstrate how the second condition for the feasibility of loop invariants is checked using the following loop invariant candidate P_{inv} as an example:

$$\begin{aligned} \{ (R_L \mid r_R), U_{TL}, P_L, D_{TL}, (X_L \mid x_R) \} &:= \text{CG}(A, R_L e_0, X_L e_0) \\ u_{TR} &:= - \left(P_L^T A P_L \right)^{-1} P_L^T A r_R \end{aligned}$$

The loop guard G is $\|r_R\| \geq \varepsilon$, so its negation $\neg G$ is $\|r_R\| < \varepsilon$. Now, it has to be checked whether P_{inv} and $\neg G$ imply the postcondition P_{post} :

$$\begin{aligned} A P D &= R (\underline{I} - \underline{J}) \\ P (I - U) &= \underline{R} \\ P D &= X (\underline{I} - \underline{J}) \\ \|R e_r^T\| &< \varepsilon \end{aligned}$$

To do that, we rewrite P_{inv} as

$$\begin{aligned} AP_L D_{\text{TL}} &= (R_L \mid r_R) \left(\frac{I - J \mid 0}{-e_r^T \mid 1} \right) \\ P_L (I - U_{\text{TL}}) &= R_L \\ P_L D_{\text{TL}} &= (X_L \mid x_R) \left(\frac{I - J \mid 0}{-e_r^T \mid 1} \right) \\ u_{\text{TR}} &= - \left(P_L^T A P_L \right)^{-1} P_L^T A r_R. \end{aligned}$$

Clearly, by rewriting $(R_L \mid r_R)$ as R , P_L as P and so on, one can see that the first three equations above are the same as the equations in the postcondition. Furthermore, the negation of the loop guard, $\|r_R\| < \varepsilon$, refers to the last column of $(R_L \mid r_R)$, just like $\|Re_r^T\| < \varepsilon$ refers to the last column of R . Thus, $P_{\text{inv}} \wedge \neg G$ implies the postcondition. While $P_{\text{inv}} \wedge \neg G$ also implies the equation

$$u_{\text{TR}} = - \left(P_L^T A P_L \right)^{-1} P_L^T A r_R,$$

this equation is not needed to render the postcondition true. In the algorithm, u_{TR} will be discarded.

Note that the initial partitioning exposed one additional column of each operand. R for example was partitioned into $(R_L \mid r_R \mid r_+)$. In the postcondition, R just consists of R_L and r_R . As mentioned before, this is possible because the number of columns of R is variable.

4.1.3. Algorithm Construction

The algorithm is constructed in the third step. In addition to the update, preprocessing operations have to be determined. Furthermore, in line with different rules for the initial partitioning, the repartitioning is modified. The process of identifying the update operations does not change.

Preprocessing

As mentioned before, the preprocessing operations are obtained by applying the initial partitioning (Table 4.1) to all but the initial node contained in the loop invariant. The first node is excluded because when the initial partitioning is applied to it, it always reduces to an expression similar to $k_R := KS(A, k_R)$ for the Krylov sequence. Since the output of this function is already known, nothing has to be computed. A FLAME worksheet, extended by the preprocessing, is shown in Figure 4.1.

4. Derivation of Algorithms for Iterative Methods

Algorithm: ...
{ P _{pre} }
Partition
Preprocessing
{ P _{inv} }
While G do
{ (P _{inv}) ∧ (G) }
Repartition
{ P _{before} }
Update
{ P _{after} }
Continue with
{ P _{inv} }
endwhile
{ (P _{inv}) ∧ ¬ (G) }
{ P _{post} }

Figure 4.1.: The skeleton of a FLAME worksheet, extended by some preprocessing.

Repartitioning the Operands

To ensure that the resulting algorithms make progress, the operands have to be repartitioned. The sizes of some operands depend on the number of iterations that is computed, so with every iteration, their sizes have to grow. This is done by adding rows and/or columns in the “Continue with” repartitioning. The rules are shown in Table 4.3.

Initial Partitioning	Repartition	Continue with
$(B_L b_R)$	$(B_0 b_1)$	$(B_0 b_1 b_2)$
$\left(\frac{B_T}{b_R} \right)$	$\left(\frac{B_0}{b_1} \right)$	$\left(\frac{B_0}{b_1} \right)$
$\left(\frac{B_{TL} b_{TR}}{b_{BL} \beta_{BR}} \right)$	$\left(\frac{B_{00} b_{01}}{b_{10} \beta_{11}} \right)$	$\left(\frac{B_{00} b_{01} b_{02}}{b_{10} \beta_{11} \beta_{12}} \right)$
$(B_L b_R b_+)$	$(B_0 b_1 b_2)$	$(B_0 b_1 b_2 b_3)$

Table 4.3.: “Repartition” and “Continue with” rules for iterative methods.

Predicates P_{before} and P_{after}

To obtain the predicates P_{before} and P_{after} , the repartitioned operands are plugged into the loop invariant. The resulting expressions are flattened, using the PME if necessary. What is obtained by applying the “Repartition” rules to the loop invariant becomes P_{before} . Applying the “Continue with” partitioning result in P_{after} .

We demonstrate this for the loop invariant

$$\left((K_L \mid k_R) := KS(A, K_L e_0) \mid \neq \right).$$

The “Repartition” and “Continue with” rules for K are shown below, introducing the newly added k_3 :

$$\begin{aligned} (K_L \mid k_R \mid k_+) &\rightarrow (K_0 \mid k_1 \mid k_2) \\ (K_L \mid k_R \mid k_+) &\leftarrow (K_0 \mid k_1 \mid k_2 \mid k_3). \end{aligned}$$

Applying the “Repartition” rules to the loop invariant yields the following predicate P_{before} :

$$(K_0 \mid k_1) := KS(A, K_0 e_0)$$

Using the “Continue with” repartitioning, we obtain the expression

$$(K_0 \mid k_1 \mid k_2) := KS(A, K_0 e_0).$$

To flatten this expression, the PME is used. This results in the following P_{after} :

$$\left((K_0 \mid k_1) := KS(A, K_0 e_0) \mid k_2 := Ak_1 \right)$$

Finding the Updates

The difference between P_{after} and P_{before} now is the update. Identifying the differences is much easier compared to direct methods because they are always entire equations, not subexpression. The reason is that no quantities are updated.

For the Krylov sequence, $k_2 := Ak_1$ can easily be identified as the difference between P_{after} and P_{before} .

4.1.4. Refinement

In practice, algorithms would not, and depending on the language, can not be implemented exactly like they are derived with the presented approach. Consider the

4. Derivation of Algorithms for Iterative Methods

following three assignments as an example. They are part of the update of one non-symmetric CG algorithm.

$$\begin{aligned} r_2 &:= r_1 - Ap_1\delta_{11} \\ u_{02} &:= \left(-P_0^TAP_0\right)^{-1} \cdot P_0^TAr_2 \\ v_{12} &:= -\frac{p_1^TAr_2 + p_1^TAP_0u_{02}}{p_1^TAp_1} \end{aligned}$$

To translate the assignments above to C for example, they have to be decomposed into basic operations that are implemented in a library like BLAS. Since such libraries usually do not include functions for general products of more than two quantities, auxiliary variables have to be introduced.

While the assignments could immediately be translated into Matlab code, this would not result in efficient code. Clearly, some subexpressions appear multiple times, so it is preferable to introduce auxiliary variables for those and compute their values just once. This is referred to as common subexpression elimination. While this concept is well known in the domain of compiler construction [10], to the author's knowledge there is no research on the elimination of *overlapping* common subexpressions. Overlapping common subexpression are common subexpression that can not be eliminated at the same time. Consider the three terms $P_0^TAP_0$, $p_1^TAr_2$ and $p_1^TAP_0$ for a simple example. $P_0^TAP_0$ and $p_1^TAP_0$ have AP_0 in common, $p_1^TAr_2$ and $p_1^TAP_0$ share p_1^TA . In $p_1^TAP_0$, those two common subexpression overlap since A is part of both.

Finding a replacement of subexpressions that is optimal in the sense that it has the lowest computational cost is not trivial. Inspecting the problem, one observes that a simplified version of it can be mapped to a *maximum weight matching* problem, which is solvable in polynomial time [3]. Again, consider the terms $P_0^TAP_0$, $p_1^TAr_2$ and $p_1^TAP_0$ as an example. Every expression becomes a node in a graph. Every possible replacement of a common subexpression is represented by an edge. The resulting graph is shown in Figure 4.2. The weight of each edge would be the computational

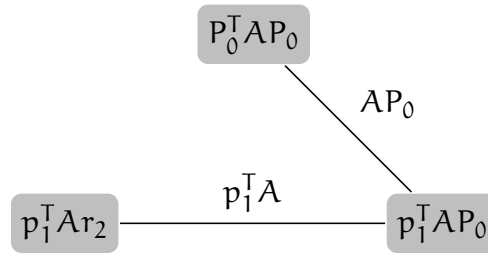


Figure 4.2.: Elimination of common subexpressions represented as a graph.

cost of the expression that is replaced. The problem then is to find a set of edges, such that each node is attached to at most one of the selected edges. At the same time, the

sum of the weights should be minimized. The requirement that each node is attached to at most one of the selected edges represents the fact that multiple replacements are not possible because the expressions overlap.

Clearly, the actual problem is more complex. A common subexpression might be replaced in more than two expressions. A graph representing this is a hypergraph. In addition to that, in longer expressions, some subexpressions do not overlap, so they can both be replaced.

In practice, trying to solve this problem is probably not necessary. The expressions encountered in most iterative methods are rarely as complex as in the example above. Furthermore, simple heuristics may already produce good results. In products of more than two quantities, matrix-vector products should always be computed first. Then, if one subexpression is replaced, all other occurrences of this expression should be replaced as well.

Alternatively, to generate as many variants as possible, all possible replacements could be constructed.

4.2. Example: Nonsymmetric CG

In this section, as a more elaborate example, we will show how the derivation of an algorithm for nonsymmetric CG proceeds.

PME Generation

The derivation of some properties for nonsymmetric CG was already shown in Section 3.2.5 and will not be repeated here. The postcondition is shown below:

$$\begin{aligned} APD &= R (\underline{I} - \underline{J}) \\ P (I - U) &= \underline{R} \\ PD &= X (\underline{I} - \underline{J}) \\ \|\text{Re}_r^T\| &< \varepsilon \end{aligned}$$

The initial partitioning is determined based on the properties of the operands. It is applied to the postcondition as well as all derived properties. The expression that is obtained from the postcondition is flattened, yielding

$$\begin{aligned} & (AP_L D_{TL} = R_L (I - J) - r_R e_r^T \mid A p_R \delta_{BR} = r_R - r_+) \\ & (P_L (I - U_{TL}) = R_L \mid -P_L u_{TR} + p_R = r_R) \\ & (P_L D_{TL} = X_L (I - J) - x_R e_r^T \mid p_R \delta_{BR} = x_R - x_+) . \end{aligned}$$

The left-hand side is now matched by the CG function, so

$$\{ (R_L \mid r_R), U_{TL}, P_L, D_{TL}, (X_L \mid x_R) \} := \text{CG} (A, R_L e_0, X_L e_0)$$

4. Derivation of Algorithms for Iterative Methods

is obtained. Thus, all quantities on the left-hand side of that assignment are considered to be known. To find assignments for the remaining unknown quantities, the three equations on the right have to be solved. One of the derived properties is that $P^T A P$ is lower triangular, so $P_L^T A P_L$ is lower triangular as well and $P_L^T A p_R$ is zero. The system would recognize that both P_L and p_R appear in $-P_L u_{TR} + p_R = r_R$ on the left-hand side of products. Hence, $P_L^T A$ is multiplied from the left to both sides of the equation to recreate those properties. The resulting equation $-P_L^T A P_L u_{TR} = P_L^T A r_R$ contains only one unknown quantity, so it is solvable. Since $P_L^T A P_L$ is lower triangular, a triangular system is identified, which can also be written as

$$u_{TR} := - \left(P_L^T A P_L \right)^{-1} P_L^T A r_R.$$

Having found an assignment for u_{TR} , it is considered known as well. By instead using that $R^T A P$ is lower triangular, a different equation for u_{TR} would have been found, resulting in a different algorithm. In practice, two separate derivation processes would be executed for both variants; here we continue just with the first one.

Now, there is only one unknown quantity left in $-P_L u_{TR} + p_R = r_R$, so the following formula is determined for p_R :

$$p_R := r_R + P_L u_{TR}$$

Because $r_R^T r_+$ is zero, r_R^T is multiplied from the left to both sides of $A p_R \delta_{BR} = r_R - r_+$. There is only one unknown in the resulting equation $r_R^T A p_R \delta_{BR} = r_R^T r_R$, so another assignment is found:

$$\delta_{BR} := \frac{r_R^T r_R}{r_R^T A p_R}$$

Alternatively, the fact that $P^T R$ is lower triangular and rectangular could be used. Finally, $A p_R \delta_{BR} = r_R - r_+$ and $p_R \delta_{BR} = x_R - x_+$ can be solved to r_+ and x_+ , respectively, completing the PME:

$$\begin{aligned} \{ (R_L \mid r_R), u_{TL}, P_L, D_{TL}, (X_L \mid x_R) \} &:= CG(A, R_L e_0, X_L e_0) \\ u_{TR} &:= - \left(P_L^T A P_L \right)^{-1} P_L^T A r_R \\ p_R &:= r_R + P_L u_{TR} \\ \delta_{BR} &:= \frac{r_R^T r_R}{r_R^T A p_R} \\ r_+ &:= r_R - A p_R \delta_{BR} \\ x_+ &:= x_R - p_R \delta_{BR} \end{aligned}$$

Loop Invariant Identification

Based on this PME, the dependency graph is constructed. Since there are six assignments in the PME, there are six nodes:

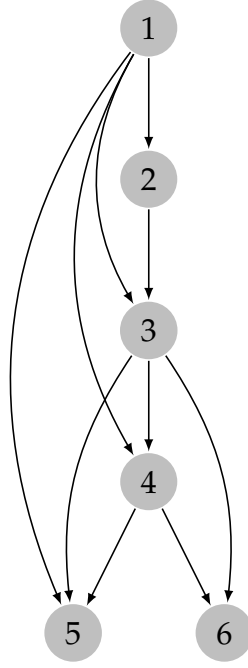


Figure 4.3.: Dependency graph for nonsymmetric CG.

1. $\{ (\begin{smallmatrix} R_L \\ r_R \end{smallmatrix}), u_{TL}, p_L, d_{TL}, (\begin{smallmatrix} x_L \\ x_R \end{smallmatrix}) \} := CG(A, R_L e_0, x_L e_0)$
2. $u_{TR} := - (p_L^T A p_L)^{-1} p_L^T A r_R$
3. $p_R := r_R + p_L u_{TR}$
4. $\delta_{BR} := \frac{r_R^T r_R}{r_R^T A p_R}$
5. $r_+ := r_R - A p_R \delta_{BR}$
6. $x_+ := x_R - p_R \delta_{BR}$

The corresponding dependency graph is shown in Figure 4.3. All nonempty subsets of this graph that respect the dependencies are feasible loop invariants. Thus, the following seven loop invariants are obtained:

$$\{1\}, \{1, 2\}, \{1, 2, 3\}, \{1, 2, 3, 4\}, \{1, 2, 3, 4, 5\}, \{1, 2, 3, 4, 6\}, \{1, 2, 3, 4, 5, 6\}$$

To determine the loop guard G , the additional predicate in the postcondition is inspected. It is $\|Re_r^T\| < \varepsilon$. Since the precondition contains `FirstColumnInput[R]`, according to Table 4.2, the loop guard $\|r_R\| \geq \varepsilon$ is selected.

4. Derivation of Algorithms for Iterative Methods

In this example, we show the derivation for the set $\{1, 2, 3\}$, so the loop invariant is

$$\begin{aligned} P_{\text{inv}} = \{ \{ (R_L \mid r_R), U_{TL}, P_L, D_{TL}, (X_L \mid x_R) \} &:= CG(A, R_L e_0, X_L e_0) \wedge \\ u_{TR} &:= - (P_L^T A P_L)^{-1} P_L^T A r_R \wedge \\ p_R &:= r_R + P_L u_{TR} \}. \end{aligned}$$

Algorithm Construction

In a first step, the preprocessing operations are determined. The relevant assignments are the ones contained in the loop invariant, except for the one obtained from the first node. They are shown below:

$$\begin{aligned} u_{TR} &:= - (P_L^T A P_L)^{-1} P_L^T A r_R \\ p_R &:= r_R + P_L u_{TR} \end{aligned}$$

According to the initial partitioning, P_L is empty, that is, it has the size $n \times 0$. Consequently, the right-hand side of the first assignment is empty, too, so it disappears. The second assignment reduces to

$$p_R := r_R,$$

which is the only preprocessing operation. Next, the update is derived. The “Repartition” rules are determined using Table 4.3:

$$\begin{aligned} (R_L \mid r_R \mid r_+) &\rightarrow (R_0 \mid r_1 \mid r_2) & (X_L \mid x_R \mid x_+) &\rightarrow (X_0 \mid x_1 \mid x_2) \\ \left(\begin{array}{c|c} U_{TL} & u_{TR} \\ \hline 0 & 0 \end{array} \right) &\rightarrow \left(\begin{array}{c|c} U_{00} & u_{01} \\ \hline 0 & 0 \end{array} \right) & \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & \delta_{BR} \end{array} \right) &\rightarrow \left(\begin{array}{c|c} D_{00} & 0 \\ \hline 0 & \delta_{11} \end{array} \right) \\ (P_L \mid p_R) &\rightarrow (P_0 \mid p_1) \end{aligned}$$

Applying those rules to the loop invariant yields the following predicate P_{before} :

$$\begin{aligned} P_{\text{before}} = \{ \{ (R_0 \mid r_1), U_{00}, P_0, D_{00}, (X_0 \mid x_1) \} &:= CG(A, R_0 e_0, X_0 e_0) \wedge \\ u_{01} &:= - (P_0^T A P_0)^{-1} P_0^T A r_1 \wedge \\ p_1 &:= r_1 + P_0 u_{01} \} \end{aligned}$$

The “Continue with” repartitioning is

$$\begin{aligned}
 (R_L \mid r_R \mid r_+) &\leftarrow (R_0 \mid r_1 \mid r_2 \mid r_3) \\
 (X_L \mid x_R \mid x_+) &\leftarrow (X_0 \mid x_1 \mid x_2 \mid x_3) \\
 \left(\frac{U_{TL}}{0} \mid \frac{u_{TR}}{0} \right) &\leftarrow \left(\frac{U_{00} \mid u_{01} \mid u_{02}}{0 \mid 0 \mid 0} \right) \\
 \left(\frac{D_{TL}}{0} \mid \frac{0}{\delta_{BR}} \right) &\leftarrow \left(\frac{D_{00} \mid 0 \mid 0}{0 \mid \delta_{11} \mid 0} \right) \\
 (P_L \mid p_R) &\leftarrow (P_0 \mid p_1 \mid p_2).
 \end{aligned}$$

Plugging that into the function, the following expression is obtained:

$$\left\{ (R_0 \mid r_1 \mid r_2), \left(\frac{U_{00} \mid u_{01}}{0 \mid 0} \right), (P_0 \mid p_1), \right. \\
 \left. \left(\frac{D_{00} \mid 0}{0 \mid \delta_{11}} \right), (X_0 \mid x_1 \mid x_2) \right\} := \text{CG}(A, R_0 e_0, X_0 e_0)$$

It is flattened by using the PME, resulting in six assignments:

$$\begin{aligned}
 \{ (R_0 \mid r_1), U_{00}, P_0, D_{00}, (X_0 \mid x_1) \} &:= \text{CG}(A, R_0 e_0, X_0 e_0) \\
 u_{01} &:= - (P_0^T A P_0)^{-1} P_0^T A r_1 \\
 p_1 &:= r_1 + P_0 u_{01} \\
 \delta_{11} &:= \frac{r_1^T r_1}{r_1^T A p_1} \\
 r_2 &:= r_1 - A p_1 \delta_{11} \\
 x_2 &:= x_1 - p_1 \delta_{11}
 \end{aligned}$$

For the second assignment of the loop invariant, the PME of a lower triangular system is needed (see Section 2.1).

$$\begin{aligned}
 \left(\frac{u_{02}}{v_{12}} \right) &:= - \left(\frac{P_0^T A P_0 \mid 0}{p_1^T A P_0 \mid p_1^T A p_1} \right)^{-1} \left(\frac{P_0^T A r_2}{p_1^T A r_2} \right) \\
 \Rightarrow \begin{cases} u_{02} &:= (-P_0^T A P_0)^{-1} \cdot P_0^T A r_2 \\ v_{12} &:= -\frac{p_1^T A r_2 + p_1^T A P_0 u_{02}}{p_1^T A p_1} \end{cases}
 \end{aligned}$$

4. Derivation of Algorithms for Iterative Methods

The third equation, $p_1 := r_1 + P_0 u_{01}$, becomes $p_2 := r_2 + P_0 u_{02} + p_1 v_{12}$ after the application of the “Continue with” partitioning. Now that the complete predicate P_{after} is determined, the update is found by comparing it to P_{before} . The assignments that are contained in P_{after} , but not in P_{before} , constitute the update. They are shown below.

$$\begin{aligned}\delta_{11} &:= \frac{r_1^T r_1}{r_1^T A p_1} \\ r_2 &:= r_1 - A p_1 \delta_{11} \\ x_2 &:= x_1 - p_1 \delta_{11} \\ u_{02} &:= \left(-P_0^T A P_0\right)^{-1} \cdot P_0^T A r_2 \\ v_{12} &:= -\frac{p_1^T A r_2 + p_1^T A P_0 u_{02}}{p_1^T A p_1} \\ p_2 &:= r_2 + P_0 u_{02} + p_1 v_{12}\end{aligned}$$

Postprocessing

For this example, we use the heuristics explained in Section 4.1.4 for the elimination of common subexpressions. The first expression that is identified is $A p_1$. The auxiliary variable $t_1 := A p_1$ is introduced and all occurrences of $A p_1$ are replaced with t_1 . The resulting assignments are shown below.

$$\begin{aligned}t_1 &:= A p_1 \\ \delta_{11} &:= \frac{r_1^T r_1}{r_1^T t_1} \\ r_2 &:= r_1 - t_1 \delta_{11} \\ x_2 &:= x_1 - p_1 \delta_{11} \\ u_{02} &:= \left(-P_0^T A P_0\right)^{-1} \cdot P_0^T A r_2 \\ v_{12} &:= -\frac{p_1^T A r_2 + p_1^T A P_0 u_{02}}{p_1^T t_1} \\ p_2 &:= r_2 + P_0 u_{02} + p_1 v_{12}\end{aligned}$$

Further auxiliary variables are introduced for $A r_2$ and $P_0 u_{02}$. The update that is obtained at the end of this step is shown in the filled out worksheet in Figure 4.4. Some of the partitionings and the loop invariant are omitted in the interest of legibility.

4.3. Example: BiCG

As a second example, we show the derivation of two algorithms for the biconjugate gradient method (BiCG). Since we already showed a full example in the previous

Algorithm: nonsymmetric CG
Partition $R \rightarrow (R_L \mid r_R \mid r_+)$ where R_L is $n \times 0$ $p_R := r_R$ While $\ r_R\ \geq \varepsilon$ do Repartition $(R_L \mid r_R \mid r_+) \rightarrow (R_0 \mid r_1 \mid r_2)$ <hr/> $t_1 := Ap_1$ $\delta_{11} := \frac{r_1^T r_1}{r_1^T t_1}$ $r_2 := r_1 - t_1 \delta_{11}$ $x_2 := x_1 - p_1 \delta_{11}$ $t_2 := Ar_2$ $u_{02} := (-P_0^T A P_0)^{-1} \cdot P_0^T t_2$ $t_3 := P_0 u_{02}$ $v_{12} := -\frac{p_1^T t_2 + p_1^T A t_3}{p_1^T t_1}$ $p_2 := r_2 + t_3 + p_1 v_{12}$ <hr/> Continue with $(R_L \mid r_R \mid r_+) \leftarrow (R_0 \mid r_1 \mid r_2 \mid r_3)$ endwhile

Figure 4.4.: Worksheet for a nonsymmetric CG algorithm.

4. Derivation of Algorithms for Iterative Methods

section, we now proceed at a slightly higher pace. Pre- and postcondition are shown below.

$$\begin{aligned}
 P_{\text{pre}} : \{ & \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\
 & \text{Output}[P] \wedge \text{Matrix}[P] \wedge \\
 & \text{Output}[\tilde{P}] \wedge \text{Matrix}[\tilde{P}] \wedge \\
 & \text{Output}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge \\
 & \text{FirstColumnInput}[R] \wedge \text{Matrix}[R] \wedge \\
 & \text{FirstColumnInput}[\underline{R}] \wedge \text{Matrix}[\underline{R}] \wedge \\
 & \text{FirstColumnInput}[\tilde{R}] \wedge \text{Matrix}[\tilde{R}] \wedge \\
 & \text{FirstColumnInput}[\underline{\tilde{R}}] \wedge \text{Matrix}[\underline{\tilde{R}}] \wedge \\
 & \text{Diagonal}[R^T \tilde{R}] \wedge \text{Diagonal}[\underline{R}^T \underline{\tilde{R}}] \wedge \\
 & \text{Diagonal}[\underline{R}^T \tilde{R}] \wedge \text{Diagonal}[\underline{R}^T \underline{\tilde{R}}] \wedge \\
 & \text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \wedge \\
 & \text{Output}[U] \wedge \text{Matrix}[U] \wedge \text{UpperDiagonal}[U] \wedge \\
 & \text{Matrix}[\underline{I} - \underline{J}] \wedge \text{LowerTrapezoidal}[\underline{I} - \underline{J}] \}
 \end{aligned}$$

$$\begin{aligned}
 P_{\text{post}} : \{ & APD = R (\underline{I} - \underline{J}) \\
 & A^T \tilde{P} D = \tilde{R} (\underline{I} - \underline{J}) \\
 & P (I - U) = \underline{R} \\
 & \tilde{P} (I - U) = \underline{\tilde{R}} \\
 & PD = X (\underline{I} - \underline{J}) \\
 & \|\text{Re}_r^T\| < \varepsilon \}
 \end{aligned}$$

Following from the precondition, the function representing the operation is

$$\{R, \tilde{R}, U, P, \tilde{P}, D, X\} := \text{BiCG}(A, \text{Re}_0, \text{Xe}_0).$$

Derivation of Properties

For BiCG, neither R nor \tilde{R} is orthogonal. Instead, they are mutually orthogonal, which means that $R^T \tilde{R}$ is diagonal. Except for that, the derivation of properties is mostly the same to the one shown in the example in Section 3.2.5. For this reason, it will not be shown here. Instead, we give a short overview of the relevant properties:

- $\underline{\tilde{R}}^T AP$ and $\underline{R}^T A^T \tilde{P}$ are lower triangular.
- $\tilde{P}^T AP$ (and thus $P^T A^T \tilde{P}$) is diagonal.
- $P^T \tilde{R}$ and $\tilde{P}^T R$ are lower triangular and rectangular.

PME Generation

As a first step towards the PME, the operands are partitioned. Initially, only the first columns of R and \tilde{R} are known, so they are partitioned into $(R_L | r_R | r_+)$ and $(\tilde{R}_L | \tilde{r}_R | \tilde{r}_+)$, respectively. The remaining operands are partitioned accordingly:

$$\begin{aligned}
 A (P_L | p_R) \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & \delta_{BR} \end{array} \right) &= (R_L | r_R | r_+) \left(\begin{array}{c|c} I-J & 0 \\ \hline -e_r^T & 1 \\ \hline 0 & -1 \end{array} \right) \\
 A^T (\tilde{P}_L | \tilde{p}_R) \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & \delta_{BR} \end{array} \right) &= (\tilde{R}_L | \tilde{r}_R | \tilde{r}_+) \left(\begin{array}{c|c} I-J & 0 \\ \hline -e_r^T & 1 \\ \hline 0 & -1 \end{array} \right) \\
 (P_L | p_R) \left(\begin{array}{c|c} I-U_{TL} & -u_{TR} \\ \hline 0 & 1 \end{array} \right) &= (R_L | r_R) \\
 (\tilde{P}_L | \tilde{p}_R) \left(\begin{array}{c|c} I-U_{TL} & -u_{TR} \\ \hline 0 & 1 \end{array} \right) &= (\tilde{R}_L | \tilde{r}_R) \\
 (P_L | p_R) \left(\begin{array}{c|c} D_{TL} & 0 \\ \hline 0 & \delta_{BR} \end{array} \right) &= (X_L | x_R | x_+) \left(\begin{array}{c|c} I-J & 0 \\ \hline -e_r^T & 1 \\ \hline 0 & -1 \end{array} \right)
 \end{aligned}$$

After the execution of the Matrix Arithmetic step, those expressions become

$$\begin{aligned}
 (AP_L D_{TL} &= R_L (I-J) - r_R e_r^T \mid Ap_R \delta_{BR} = r_R - r_+) \\
 (A^T \tilde{P}_L D_{TL} &= \tilde{R}_L (I-J) - \tilde{r}_R e_r^T \mid A\tilde{p}_R \delta_{BR} = \tilde{r}_R - \tilde{r}_+) \\
 (P_L (I-U_{TL}) &= R_L \mid -P_L u_{TR} + p_R = r_R) \\
 (\tilde{P}_L (I-U_{TL}) &= \tilde{R}_L \mid -\tilde{P}_L u_{TR} + \tilde{p}_R = \tilde{r}_R) \\
 (P_L D_{TL} &= X_L (I-J) - x_R e_r^T \mid p_R \delta_{BR} = x_R - x_+).
 \end{aligned}$$

The left-hand sides of those expressions can now be replaced by the BiCG function:

$$\{(R_L | r_R), (\tilde{R}_L | \tilde{r}_R), U_{TL}, P_L, \tilde{P}_L, D_{TL}, (X_L | x_R)\} := \text{BiCG}(A, R_L e_0, X_L e_0)$$

Updates for the remaining unknown quantities are obtained by solving the equations on the right-hand using the derived properties. For this example, to find an assignment for u_{TR} , we use that $\tilde{P}^T A P$ is diagonal. Multiplying $\tilde{P}_L^T A$ from the left to both sides of $-P_L u_{TR} + p_R = r_R$ results in $-\tilde{P}_L^T A P_L u_{TR} = \tilde{P}_L^T A r_R$. Since $\tilde{P}_L^T A P_L$ is diagonal, and thus is invertible, the following assignment is obtained:

$$u_{TR} := -(\tilde{P}_L^T A P_L)^{-1} \tilde{P}_L^T A r_R$$

4. Derivation of Algorithms for Iterative Methods

Note that it would have been also possible to use $-\tilde{P}_L u_{TR} + \tilde{p}_R = \tilde{r}_R$ to get to an assignment for u_{TR} . Now, the following assignments are obtained for p_R and \tilde{p}_R :

$$\begin{aligned} p_R &:= r_R + P_L u_{TR} \\ \tilde{p}_R &:= \tilde{r}_R + \tilde{P}_L u_{TR} \end{aligned}$$

Similarly to u_{TR} , δ_{BR} can be computed in several different ways. Here, we use that $\tilde{P}^T R$ is lower triangular and rectangular, in combination with the equation $A p_R \delta_{BR} = r_R - r_+$. Multiplying \tilde{p}_R^T from the left and solving to δ_{BR} yields

$$\delta_{BR} := \frac{\tilde{p}_R^T r_R}{\tilde{p}_R^T A p_R}$$

Finally, the following assignments are derived for r_+ , \tilde{r}_+ and x_+ :

$$\begin{aligned} r_+ &:= r_R - A p_R \delta_{BR} \\ \tilde{r}_+ &:= \tilde{r}_R - A \tilde{p}_R \delta_{BR} \\ x_+ &:= x_R - p_R \delta_{BR} \end{aligned}$$

The complete PME is shown below (already in form of a list, in anticipation of the construction of the dependency graph):

1. $\{ (R_L \mid r_R), (\tilde{R}_L \mid \tilde{r}_R), u_{TL}, P_L, \tilde{P}_L, D_{TL}, (X_L \mid x_R) \} := \text{BiCG}(A, R_L e_0, X_L e_0)$
2. $u_{TR} := -(\tilde{P}_L^T A P_L)^{-1} \tilde{P}_L^T A r_R$
3. $p_R := r_R + P_L u_{TR}$
4. $\tilde{p}_R := \tilde{r}_R + \tilde{P}_L u_{TR}$
5. $\delta_{BR} := \frac{\tilde{p}_R^T r_R}{\tilde{p}_R^T A p_R}$
6. $r_+ := r_R - A p_R \delta_{BR}$
7. $\tilde{r}_+ := \tilde{r}_R - A \tilde{p}_R \delta_{BR}$
8. $x_+ := x_R - p_R \delta_{BR}$

Loop Invariant Identification

The dependency graph is shown in Figure 4.5. Of all the subsets of this graph that respect the dependencies, only the empty one fails to satisfy the conditions for the feasibility of loop invariant. The remaining 13 subsets are feasible loop invariants. We will continue the derivation in this example with the loop invariant that corresponds to the full set.

Because of the additional predicate in the postcondition, $\|Re_r^T\| < \varepsilon$, and the property $\text{FirstColumnInput}[R]$ in the precondition, the loop guard $\|r_R\| < \varepsilon$ is determined.

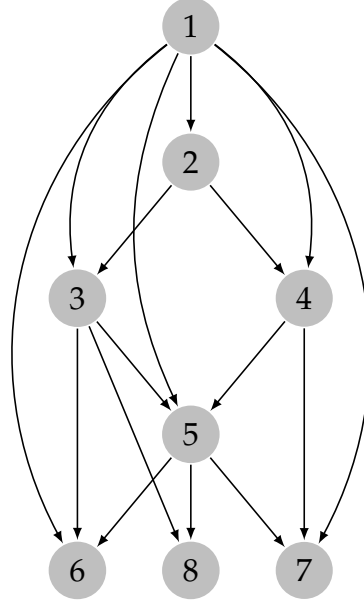


Figure 4.5.: Dependency graph for BiCG.

Algorithm Construction

The first part of this step consists of finding the preprocessing operations. This is done by taking all assignments of the loop invariant except for the one from the first node, and eliminating all expression that are empty in the initial partitioning. The relevant assignments are shown below:

$$u_{TR} := - \left(\tilde{P}_L^T A P_L \right)^{-1} \tilde{P}_L^T A r_R$$

$$p_R := r_R + P_L u_{TR}$$

$$\tilde{p}_R := \tilde{r}_R + \tilde{P}_L u_{TR}$$

$$\delta_{BR} := \frac{\tilde{p}_R^T r_R}{\tilde{p}_R^T A p_R}$$

$$r_+ := r_R - A p_R \delta_{BR}$$

$$\tilde{r}_+ := \tilde{r}_R - A \tilde{p}_R \delta_{BR}$$

$$x_+ := x_R - p_R \delta_{BR}$$

Initially, P_L and \tilde{P}_L have the size $n \times 0$ and u_{TR} is of size 0×1 . As a result, the first assignment is eliminated, while the second and third are reduced to

$$p_R := r_R$$

$$\tilde{p}_R := \tilde{r}_R.$$

4. Derivation of Algorithms for Iterative Methods

The remaining four assignments do not change.

To determine the update, the predicates P_{before} and P_{after} have to be constructed by repartitioning the loop invariant. The “Repartition” rules are shown below:

$$\begin{aligned}
 (R_L \mid r_R \mid r_+) &\rightarrow (R_0 \mid r_1 \mid r_2) & (\tilde{R}_L \mid \tilde{r}_R \mid \tilde{r}_+) &\rightarrow (\tilde{R}_0 \mid \tilde{r}_1 \mid \tilde{r}_2) \\
 \left(\frac{U_{TL}}{0} \mid \frac{u_{TR}}{0} \right) &\rightarrow \left(\frac{U_{00}}{0} \mid \frac{u_{01}}{0} \right) & \left(\frac{D_{TL}}{0} \mid \frac{0}{\delta_{BR}} \right) &\rightarrow \left(\frac{D_{00}}{0} \mid \frac{0}{\delta_{11}} \right) \\
 (P_L \mid p_R) &\rightarrow (P_0 \mid p_1) & (\tilde{P}_L \mid \tilde{p}_R) &\rightarrow (\tilde{P}_0 \mid \tilde{p}_1) \\
 (X_L \mid x_R \mid x_+) &\rightarrow (X_0 \mid x_1 \mid x_2)
 \end{aligned}$$

Applying that to the loop invariant yields the following equations which constitute P_{before} :

$$\begin{aligned}
 \{ (R_0 \mid r_1), (\tilde{R}_0 \mid \tilde{r}_1), U_{00}, P_0, \tilde{P}_0, D_{00}, (X_0 \mid x_1) \} &:= \text{BiCG}(A, R_0 e_0, X_0 e_0) \\
 u_{01} &:= -(\tilde{P}_0^T A P_0)^{-1} \tilde{P}_0^T A r_1 \\
 p_1 &:= r_1 + P_0 u_{01} \\
 \tilde{p}_1 &:= \tilde{r}_1 + \tilde{P}_0 u_{01} \\
 \delta_{11} &:= \frac{\tilde{p}_1^T r_1}{\tilde{p}_1^T A p_1} \\
 r_2 &:= r_1 - A p_1 \delta_{11} \\
 \tilde{r}_2 &:= \tilde{r}_1 - A \tilde{p}_1 \delta_{11} \\
 x_2 &:= x_1 - p_1 \delta_{11}
 \end{aligned}$$

For the “Continue with” repartitioning it is important to note that U is upper diagonal. Thus, u_{02} is zero:

$$\begin{aligned}
 (R_L \mid r_R \mid r_+) &\leftarrow (R_0 \mid r_1 \mid r_2 \mid r_3) \\
 (\tilde{R}_L \mid \tilde{r}_R \mid \tilde{r}_+) &\leftarrow (\tilde{R}_0 \mid \tilde{r}_1 \mid \tilde{r}_2 \mid \tilde{r}_3) \\
 \left(\frac{U_{TL}}{0} \mid \frac{u_{TR}}{0} \right) &\leftarrow \left(\frac{U_{00}}{0} \mid \frac{u_{01}}{0} \mid \frac{0}{\gamma_{12}} \right) \\
 \left(\frac{D_{TL}}{0} \mid \frac{0}{\delta_{BR}} \right) &\leftarrow \left(\frac{D_{00}}{0} \mid \frac{0}{\delta_{11}} \mid \frac{0}{\delta_{22}} \right) \\
 (P_L \mid p_R) &\leftarrow (P_0 \mid p_1 \mid p_2) \\
 (\tilde{P}_L \mid \tilde{p}_R) &\leftarrow (\tilde{P}_0 \mid \tilde{p}_1 \mid \tilde{p}_2) \\
 (X_L \mid x_R \mid x_+) &\leftarrow (X_0 \mid x_1 \mid x_2 \mid x_3)
 \end{aligned}$$

This repartitioning transforms

$$\{ (\mathbf{R}_L \mid \mathbf{r}_R), (\tilde{\mathbf{R}}_L \mid \tilde{\mathbf{r}}_R), \mathbf{u}_{TL}, \mathbf{p}_L, \tilde{\mathbf{p}}_L, \mathbf{D}_{TL}, (\mathbf{x}_L \mid \mathbf{x}_R) \} := \text{BiCG} (A, \mathbf{R}_L \mathbf{e}_0, \mathbf{x}_L \mathbf{e}_0)$$

into those assignments that are also contained in the predicate P_{before} . Applying it to

$$\mathbf{u}_{TR} := - \left(\tilde{\mathbf{p}}_L^T A \mathbf{p}_L \right)^{-1} \tilde{\mathbf{p}}_L^T A \mathbf{r}_R,$$

yields

$$\left(\begin{array}{c} 0 \\ \mathbf{v}_{12} \end{array} \right) := \left(\begin{array}{c|c} \tilde{\mathbf{p}}_0^T A \mathbf{p}_0 & 0 \\ \hline 0 & \tilde{\mathbf{p}}_1^T A \mathbf{p}_1 \end{array} \right)^{-1} \left(\begin{array}{c} \tilde{\mathbf{p}}_0^T A \mathbf{r}_2 \\ \tilde{\mathbf{p}}_1^T A \mathbf{r}_2 \end{array} \right).$$

Flattening this expression results in the assignment

$$\mathbf{v}_{12} := \frac{\tilde{\mathbf{p}}_1^T A \mathbf{r}_2}{\tilde{\mathbf{p}}_1^T A \mathbf{p}_1}$$

for \mathbf{v}_{12} . Because of

$$\begin{aligned} \mathbf{p}_R &:= \mathbf{r}_R + \mathbf{p}_L \mathbf{u}_{TR} \\ \tilde{\mathbf{p}}_R &:= \tilde{\mathbf{r}}_R + \tilde{\mathbf{p}}_L \mathbf{u}_{TR}, \end{aligned}$$

the following two expressions are obtained and added to P_{after} :

$$\begin{aligned} \mathbf{p}_2 &:= \mathbf{r}_2 + \mathbf{p}_1 \mathbf{v}_{12} \\ \tilde{\mathbf{p}}_2 &:= \tilde{\mathbf{r}}_2 + \tilde{\mathbf{p}}_1 \mathbf{v}_{12} \end{aligned}$$

For the remaining assignments of the loop invariant, only the indices change:

$$\begin{aligned} \delta_{22} &:= \frac{\tilde{\mathbf{p}}_2^T \mathbf{r}_2}{\tilde{\mathbf{p}}_2^T A \mathbf{p}_2} \\ \mathbf{r}_3 &:= \mathbf{r}_2 - A \mathbf{p}_2 \delta_{22} \\ \tilde{\mathbf{r}}_3 &:= \tilde{\mathbf{r}}_2 - A \tilde{\mathbf{p}}_2 \delta_{22} \\ \mathbf{x}_3 &:= \mathbf{x}_2 - \mathbf{p}_2 \delta_{22} \end{aligned}$$

Now that the P_{before} and P_{after} are completely determined, the update is found by identifying those assignments that are contained in P_{after} , but not in P_{before} . For this example, no common subexpressions are replaced. The worksheet for this algorithm, together with a second one for the algorithm obtained from the set that only contains the first node, is shown in Figure 4.6.

4. Derivation of Algorithms for Iterative Methods

Algorithm: BiCG	Algorithm: BiCG
Partition $R \rightarrow (R_L \mid r_R \mid r_+)$ where R_L is $n \times 0$ $p_R := r_R$ $\tilde{p}_R := \tilde{r}_R$ $\delta_{BR} := \frac{\tilde{p}_R^T r_R}{\tilde{p}_R^T A p_R}$ $r_+ := r_R - A p_R \delta_{BR}$ $\tilde{r}_+ := \tilde{r}_R - A \tilde{p}_R \delta_{BR}$ $x_+ := x_R - p_R \delta_{BR}$ While $\ r_R\ \geq \varepsilon$ do Repartition $(R_L \mid r_R \mid r_+) \rightarrow (R_0 \mid r_1 \mid r_2)$ <hr/> $v_{12} := \frac{\tilde{p}_1^T A r_2}{\tilde{p}_1^T A p_1}$ $p_2 := r_2 + p_1 v_{12}$ $\tilde{p}_2 := \tilde{r}_2 + \tilde{p}_1 v_{12}$ $\delta_{22} := \frac{\tilde{p}_2^T r_2}{\tilde{p}_2^T A p_2}$ $r_3 := r_2 - A p_2 \delta_{22}$ $\tilde{r}_3 := \tilde{r}_2 - A \tilde{p}_2 \delta_{22}$ $x_3 := x_2 - p_2 \delta_{22}$ <hr/> Continue with $(R_L \mid r_R \mid r_+) \leftarrow (R_0 \mid r_1 \mid r_2 \mid r_3)$ endwhile	Partition $R \rightarrow (R_L \mid r_R \mid r_+)$ where R_L is $n \times 0$ While $\ r_R\ \geq \varepsilon$ do Repartition $(R_L \mid r_R \mid r_+) \rightarrow (R_0 \mid r_1 \mid r_2)$ <hr/> $u_{01} := - \left(\tilde{p}_0^T A p_0 \right)^{-1} \tilde{p}_0^T A r_1$ $p_1 := r_1 + p_0 u_{01}$ $\tilde{p}_1 := \tilde{r}_1 + \tilde{p}_0 u_{01}$ $\delta_{11} := \frac{\tilde{p}_1^T r_1}{\tilde{p}_1^T A p_1}$ $r_2 := r_1 - A p_1 \delta_{11}$ $\tilde{r}_2 := \tilde{r}_1 - A \tilde{p}_1 \delta_{11}$ $x_2 := x_1 - p_1 \delta_{11}$ <hr/> Continue with $(R_L \mid r_R \mid r_+) \leftarrow (R_0 \mid r_1 \mid r_2 \mid r_3)$ endwhile

Figure 4.6.: Worksheets for two BiCG algorithms. The one on the left is obtained from the loop invariant that consists of the full set of nodes. The right one is derived from the set that only contains the first node.

4.3.1. Remark on the Equivalence of Loop Invariants and Algorithms

Comparing those algorithms, one notes that the updates are very similar. In fact, according to the criteria established in Section 2.4, those algorithms are considered equivalent. With a suitable replacement, the loop invariant for one can be transformed into the loop invariant of the other. The differences in the shape of some assignments stem from the fact that u_{01} is zero except for the last position. In the algorithm on the left in Figure 4.6, this is revealed by the repartitioning, while this is not the case on the right.

In general, one observes that the presented approach always produces two equivalent loop invariants, namely the subset of the dependency graph that only contains the first node, and the full set. For direct methods, the coarsest possible, that is, the standard 2×2 partitioning, never results in equivalent loop invariants. Finer partitionings result in new loop invariants than can not be found with a coarser one.

For the iterative methods covered in this thesis, this is different. The partitioning that is used for the derivation is the coarsest possible one, as a coarser one would only partition those matrices that are initially partially known. All other matrices would not be partitioned at all, so it would not be possible to derive a PME. On the other hand, for quite a few of the methods presented in Appendix B, a finer partitioning does not result in new loop invariants.

This is the case for those methods where either there is no matrix U or it is upper diagonal. The reason is that for those methods, finer partitionings do not partition quantities that are computed in one iteration in multiple parts. All they do is expose additional quantities that are fully computed in different iterations.

4.4. Scope and Limitations

The presented method extends to a lot more iterative methods than those shown as examples throughout the thesis. Matrix representations of further methods are shown in Appendix B. Note that this includes stationary iterative methods (B.2). For those, the derivation is even simpler because no properties have to be derived.

With the presented approach, it is possible to derive a large number of algorithms for most iterative methods. In case of CG for example, by using the derived properties to solve equations in different ways, four PMEs are found. From each PME, seven loop invariants are obtained, resulting in 28 algorithms. Another considerable factor is added by the elimination of common subexpressions. Similar numbers can be expected for other, comparably complex iterative methods.

There are, however, some limitations. They are explained in the following.

Rewriting of Updates

In some cases, the presented method is not able to generate those assignments that are commonly found in literature. This is for example the case for symmetric CG. The derived updates for v_{12} always have a shape like this:

$$v_{12} := \frac{p_1^T A r_2}{p_1^T A p_1} \quad (4.1)$$

Usually, the following formula is used [1, 11, 13]:

$$v_{12} := \frac{r_2^T r_2}{r_1^T r_1}$$

The advantage is that the matrix-vector product $p_1^T A$ is eliminated. It is obtained as follows. We begin with rewriting $r_2 = r_1 - A p_1 \delta_{11}$, which is the update for r_2 , as $A p_1 = (r_1 - r_2) \delta_{11}^{-1}$. Then, both sides of this equation are transposed, resulting in $p_1^T A = \delta_{11}^{-1} (r_1 - r_2)^T$. Now, this equation is used to replace $p_1^T A$ in the numerator of equation (4.1):

$$v_{12} := \frac{\delta_{11}^{-1} (r_1 - r_2)^T r_2}{p_1^T A p_1}$$

Because of the orthogonality of R , it simplifies to

$$v_{12} := \frac{\delta_{11}^{-1} r_2^T r_2}{p_1^T A p_1}.$$

δ_{11} is then replaced with

$$\delta_{11} = \frac{r_1^T r_1}{p_1^T A p_1},$$

finally resulting in

$$v_{12} := \frac{p_1^T A p_1}{r_1^T r_1} \cdot \frac{r_2^T r_2}{p_1^T A p_1} = \frac{r_2^T r_2}{r_1^T r_1}.$$

By itself, this transformation is not particularly difficult. Since all steps follow well defined algebraic rules, it is not even difficult to design a system that is able to perform the individual steps of this rewriting. The problem is that based on the initial equation (4.1), there is no indication that it is possible to reduce the number of matrix-vector products. Even if that is known, there is no indication which steps to perform. Thus, any system that is supposed to rewrite the assignment has to perform some sort of exhaustive search. Unfortunately, since expressions are substituted, the search space is infinite. Heuristics have to be applied to guarantee the termination of this search. However, they must not limit the capability of the system to find such rewritings.

Normalized Vectors

Some iterative methods construct a set of orthonormal vectors. With the presented approach, it is not possible to derive algorithms for those methods. Consider the Arnoldi iteration as an example. The matrix representation is shown below. It is based on the description in [11]:

$$\{Q, H\} := AI(A, Qe_0) \equiv \begin{cases} P_{\text{pre}} : \{ \text{Input}(A) \wedge \text{Matrix}(A) \\ \quad \text{FirstColumnInput}(Q) \wedge \text{Matrix}[Q] \wedge \\ \quad \text{Orthonormal}[Q] \wedge \\ \quad \text{Output}[H] \wedge \text{Matrix}[H] \wedge \text{UpperHessenberg}[H] \} \\ P_{\text{post}} : \{ A\underline{Q} = QH \\ \quad \text{size}(Q) = n \times m \} \end{cases}$$

The postcondition is repartitioned as follows:

$$A \left(\begin{array}{c|c} Q_L & q_R \end{array} \right) = \left(\begin{array}{c|c|c} Q_L & q_R & q_+ \end{array} \right) \left(\begin{array}{c|c} H_{TL} & h_{TR} \\ \hline h_{ML} & \eta_{MR} \\ \hline 0 & \eta_{BR} \end{array} \right)$$

After that, the Matrix Arithmetic step yields the following expression:

$$\left(A Q_L = Q_L H_{TL} + q_R h_{ML} \mid A q_R = Q_L h_{TR} + q_R \eta_{MR} + q_+ \eta_{BR} \right)$$

As usual, the left-hand side is matched by the function:

$$\left\{ \left(\begin{array}{c|c} Q_L & q_R \end{array} \right), \left(\begin{array}{c} H_{TL} \\ \hline h_{ML} \end{array} \right) \right\} := AI(A, Q_L e_0)$$

Using that Q is orthonormal, the following assignments are obtained for h_{TR} and η_{MR} :

$$\begin{aligned} h_{TR} &:= Q_L^T A q_R \\ \eta_{MR} &:= q_R^T A q_R \end{aligned}$$

The problem is now to compute the normalized vector q_+ and η_{BR} . From $A q_R = Q_L h_{TR} + q_R \eta_{MR} + q_+ \eta_{BR}$, the following equation can be obtained, where all quantities on the right-hand side are known:

$$q_+ \eta_{BR} = A q_R - Q_L h_{TR} - q_R \eta_{MR}$$

Since we know that q_+ is normalized, the scalar η_{BR} has to be computed as

$$\eta_{BR} := \|A q_R - Q_L h_{TR} - q_R \eta_{MR}\|. \quad (4.2)$$

4. Derivation of Algorithms for Iterative Methods

Then, q_+ is determined by the following assignment:

$$q_+ := \frac{Aq_R - Q_L h_{TR} - q_R \eta_{MR}}{\eta_{BR}}$$

With the presented method of solving equations by applying properties, the assignment (4.2) can not be obtained. To ensure that algorithms for such methods can be derived, the presented approach for solving equations must be expanded.

5. Conclusion

This thesis introduces a methodology that allows the systematic derivation of algorithms for iterative methods; the starting point for this methodology is a formal description of an iterative method in matrix form. In addition, we presented an approach for deriving properties of matrices and matrix expressions from this representation; those properties are necessary for the derivation of algorithms.

The actual derivation of algorithms consists of four major steps. First, PME's are generated by partitioning the operands of the matrix representation and applying the derived properties to solve equations. Then, from those PME's, loop invariants are obtained. In the third step, from each loop invariant, one loop-based algorithm is constructed. Finally, common subexpressions are eliminated, generating an even larger number of algorithms.

One of the most important aspects, and indispensable for the automatic generation of libraries, is that the derived algorithms are provably correct. This is ensured by constructing them around a proof of correctness, based on the loop invariants generated in the second step.

A conscious effort was made to ensure that the entire process is systematic, that is, each step is performed according to well defined rules and no guidance by a human expert is required. This allows the approach to be implemented as a tool that automatically generates algorithms based on a formal description of the operation. We consider this to be another important step towards the automatic generation of linear algebra libraries as envisioned by the founders of the FLAME project.

As for future work, there are a number of ways to build on the results of this thesis:

Implementation Executing the presented approach by hand is a laborious and thus error-prone task, not least because it was not designed to be executed by hand. To be used productively, the presented approach should be implemented as a computer program.

Stability Analysis To assess the usefulness of the derived algorithms in practice, a stability analysis is indispensable. In [2], it was shown that the FLAME methodology can be combined with a systematic stability analysis. The presented method should be extended in a similar way.

Performance Analysis While it is desirable to derive a large number of algorithms to find new, potentially faster variants, the task of identifying them should not be left to the user. Thus, similar to a systematic stability analysis, the system

5. Conclusion

should also be able to reason about the performance of the generated algorithms and select the best ones.

Matrix Representations In this thesis, only a small number of matrix representations for iterative methods is presented. Clearly, it would be desirable to find representations of many more methods. Additionally, it might be interesting to find out if this representation reveals new insights about different iterative methods and their relations to each other.

A. Matrix Properties

In the following, we define those matrix properties used throughout the thesis that are not self-explanatory. Let $A \in \mathbb{R}^{n \times m}$ be a matrix. The elements of this matrix are denoted as a_{ij} with $i \in \{0, \dots, n-1\}$ and $j \in \{0, \dots, m-1\}$.

- Upper diagonal (UpperDiagonal): $a_{ij} = 0$ for $i+1 \neq j$ with $n = m$. Consider the matrix below as an example.

$$\begin{pmatrix} 0 & a_{01} & 0 & 0 \\ 0 & 0 & a_{12} & 0 \\ 0 & 0 & 0 & a_{23} \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

- Lower diagonal (LowerDiagonal): $a_{ij} = 0$ for $i-1 \neq j$ with $n = m$.
- Diagonal and rectangular (DiagonalR): $a_{ij} = 0$ for $i \neq j$ with $n \neq m$.
- Upper diagonal and rectangular (UpperDiagonalR): $a_{ij} = 0$ for $i+1 \neq j$ with $n \neq m$.
- Lower diagonal and rectangular (LowerDiagonalR): $a_{ij} = 0$ for $i-1 \neq j$ with $n \neq m$.
- Upper trapezoidal (UpperTrapezoidal): $a_{ij} = 0$ for $i > j$ with $n < m$. Thus, the following matrix is upper trapezoidal:

$$\begin{pmatrix} a_{00} & a_{01} & a_{02} & a_{03} \\ 0 & a_{11} & a_{12} & a_{13} \\ 0 & 0 & a_{22} & a_{23} \end{pmatrix}$$

- Lower trapezoidal (LowerTrapezoidal): $a_{ij} = 0$ for $i < j$ with $n > m$.
- Upper triangular and rectangular (UpperTriangularR): $a_{ij} = 0$ for $i > j$ with $n > m$. An example of such a matrix is shown below.

$$\begin{pmatrix} a_{00} & a_{01} & a_{02} \\ 0 & a_{11} & a_{12} \\ 0 & 0 & a_{22} \\ 0 & 0 & 0 \end{pmatrix}$$

A. Matrix Properties

- Lower triangular and rectangular (LowerTriangularR): $a_{ij} = 0$ for $i < j$ with $n < m$.
- Elements on the diagonal are zero (ZeroDiagonal): $a_{ij} = 0$ for $i = j$ with $n = m$.

B. Matrix Representations of Iterative Methods

In this chapter, a collection of matrix representations of iterative methods is provided. The representations for the Krylov sequence, Steepest Descent, symmetric CG and nonsymmetric CG are modifications of the ones introduced in [5]. The differences lie in the use of the underline. The representation for BiCG is based on the one for symmetric CG.

B.1. Krylov Subspace Methods

B.1.1. Krylov Sequence

$$P_{\text{pre}} : \{ \text{Input}(\underline{A}) \wedge \text{Matrix}(\underline{A}) \wedge \\ \text{Matrix}(\underline{J}) \wedge \text{LowerDiagonal}(\underline{R}[\underline{J}]) \wedge \\ \text{FirstColumnInput}(\underline{K}) \}$$

$$P_{\text{post}} : \{ \underline{A}\underline{K} = \underline{K}\underline{J} \\ \text{size}(\underline{K}) = n \times m \}$$

B.1.2. Steepest Descent

$$P_{\text{pre}} : \{ \text{Input}[\underline{A}] \wedge \text{Matrix}[\underline{A}] \wedge \text{NonSingular}[\underline{A}] \wedge \\ \text{Output}[\underline{D}] \wedge \text{Matrix}[\underline{D}] \wedge \text{Diagonal}[\underline{D}] \wedge \\ \text{FirstColumnInput}[\underline{R}] \wedge \text{Matrix}[\underline{R}] \wedge \\ \text{ZeroDiagonal}[\underline{R}^T \underline{R}] \wedge \text{ZeroDiagonal}[\underline{R}^T \underline{R}^T] \wedge \\ \text{FirstColumnInput}[\underline{X}] \wedge \text{Matrix}[\underline{X}] \wedge \\ \text{Matrix}[\underline{I} - \underline{J}] \wedge \text{LowerTrapezoidal}[\underline{I} - \underline{J}] \}$$

$$P_{\text{post}} : \{ \underline{A}\underline{R}\underline{D} = \underline{R}(\underline{I} - \underline{J}) \\ \underline{R}\underline{D} = \underline{X}(\underline{I} - \underline{J}) \\ \|\text{Re}_r^T\| < \varepsilon \}$$

B.1.3. Conjugate Gradient (symmetric)

$P_{\text{pre}} : \{ \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \text{Symmetric}[A] \wedge$
 $\text{Output}[P] \wedge \text{Matrix}[P] \wedge$
 $\text{Output}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge$
 $\text{FirstColumnInput}[R] \wedge \text{Matrix}[R] \wedge \text{Orthogonal}[R] \wedge$
 $\text{FirstColumnInput}[\underline{R}] \wedge \text{Matrix}[\underline{R}] \wedge \text{Orthogonal}[\underline{R}] \wedge$
 $\text{Diagonal}R[R^T \underline{R}] \wedge \text{Diagonal}R[\underline{R}^T R] \wedge$
 $\text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \wedge$
 $\text{Output}[U] \wedge \text{Matrix}[U] \wedge \text{UpperDiagonal}[U] \wedge$
 $\text{Matrix}[\underline{I} - \underline{J}] \wedge \text{LowerTrapezoidal}[\underline{I} - \underline{J}] \}$

$$\begin{aligned}
 P_{\text{post}} : \{ & APD = R (\underline{I} - \underline{J}) \\
 & P (I - U) = \underline{R} \\
 & PD = X (\underline{I} - \underline{J}) \\
 & \|Re_r^T\| < \varepsilon \}
 \end{aligned}$$

B.1.4. Conjugate Gradient (nonsymmetric)

$P_{\text{pre}} : \{ \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge$
 $\text{Output}[P] \wedge \text{Matrix}[P] \wedge$
 $\text{Output}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge$
 $\text{FirstColumnInput}[R] \wedge \text{Matrix}[R] \wedge \text{Orthogonal}[R] \wedge$
 $\text{FirstColumnInput}[\underline{R}] \wedge \text{Matrix}[\underline{R}] \wedge \text{Orthogonal}[\underline{R}] \wedge$
 $\text{Diagonal}R[R^T \underline{R}] \wedge \text{Diagonal}R[\underline{R}^T R] \wedge$
 $\text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \wedge$
 $\text{Output}[U] \wedge \text{Matrix}[U] \wedge \text{StrictlyUpperTriangular}[U] \wedge$
 $\text{Matrix}[\underline{I} - \underline{J}] \wedge \text{LowerTrapezoidal}[\underline{I} - \underline{J}] \}$

$$\begin{aligned}
 P_{\text{post}} : \{ & APD = R (\underline{I} - \underline{J}) \\
 & P (I - U) = \underline{R} \\
 & PD = X (\underline{I} - \underline{J}) \\
 & \|Re_r^T\| < \varepsilon \}
 \end{aligned}$$

B.1.5. BiCG

$$\begin{aligned}
P_{\text{pre}} : \{ & \text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\
& \text{Output}[P] \wedge \text{Matrix}[P] \wedge \\
& \text{Output}[\tilde{P}] \wedge \text{Matrix}[\tilde{P}] \wedge \\
& \text{Output}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge \\
& \text{FirstColumnInput}[R] \wedge \text{Matrix}[R] \wedge \\
& \text{FirstColumnInput}[\underline{R}] \wedge \text{Matrix}[\underline{R}] \wedge \\
& \text{FirstColumnInput}[\tilde{R}] \wedge \text{Matrix}[\tilde{R}] \wedge \\
& \text{FirstColumnInput}[\tilde{\underline{R}}] \wedge \text{Matrix}[\tilde{\underline{R}}] \wedge \\
& \text{Diagonal}[R^T \tilde{R}] \wedge \text{Diagonal}[\underline{R}^T \tilde{R}] \wedge \\
& \text{Diagonal}[\underline{R}^T \tilde{R}] \wedge \text{Diagonal}[R^T \tilde{\underline{R}}] \wedge \\
& \text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \wedge \\
& \text{Output}[U] \wedge \text{Matrix}[U] \wedge \text{UpperDiagonal}[U] \wedge \\
& \text{Matrix}[\underline{I} - \underline{J}] \wedge \text{LowerTrapezoidal}[\underline{I} - \underline{J}] \}
\end{aligned}$$

$$\begin{aligned}
P_{\text{post}} : \{ & APD = R(\underline{I} - \underline{J}) & A^T \tilde{P}D = \tilde{R}(\underline{I} - \underline{J}) \\
& P(I - U) = \underline{R} & \tilde{P}(I - U) = \tilde{\underline{R}} \\
& PD = X(\underline{I} - \underline{J}) & \|\text{Re}_r^T\| < \varepsilon \}
\end{aligned}$$

B.2. Stationary Iterative Methods

Let $Ax = b$ be a linear system. e is a column vector where all entries are one. We write A as $A = D - L - U$, where D contains the entries on the main diagonal of A , L contains the entries below the main diagonal and U the ones above the main diagonal. The representations for the Gauss-Seidel, Jacobi and Successive Overrelaxation method are based on the descriptions (in indexed notation) in [1]. The one for the Richardson iteration is based on the the description (in indexed notation) in [5].

B.2.1. Gauss-Seidel Method

$$\begin{aligned}
P_{\text{pre}} : \{ & \text{Input}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge \\
& \text{Input}[L] \wedge \text{Matrix}[L] \wedge \text{LowerTriangular}[L] \wedge \\
& \text{Input}[U] \wedge \text{Matrix}[U] \wedge \text{UpperTriangular}[U] \wedge \\
& \text{Input}[b] \wedge \text{Vector}[b] \wedge \\
& \text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \}
\end{aligned}$$

B. Matrix Representations of Iterative Methods

$$P_{\text{post}} : \{(D - L)\underline{X} = \underline{U}\underline{X} + \underline{b}e^T \\ \|\underline{X}e_r^T - \underline{X}e_{r-1}^T\| < \varepsilon\}$$

B.2.2. Jacobi Method

The precondition of the Jacobi method is identical to the one for the Gauss-Seidel method.

$$P_{\text{post}} : \{D\underline{X} = (L + U)\underline{X} + \underline{b}e^T \\ \|\underline{X}e_r^T - \underline{X}e_{r-1}^T\| < \varepsilon\}$$

B.2.3. Successive Overrelaxation Method

$$P_{\text{pre}} : \{\text{Input}[D] \wedge \text{Matrix}[D] \wedge \text{Diagonal}[D] \wedge \\ \text{Input}[L] \wedge \text{Matrix}[L] \wedge \text{LowerTriangular}[L] \wedge \\ \text{Input}[U] \wedge \text{Matrix}[U] \wedge \text{UpperTriangular}[U] \wedge \\ \text{Input}[b] \wedge \text{Vector}[b] \wedge \\ \text{Input}[\omega] \wedge \text{Scalar}[\omega] \\ \text{FirstColumnInput}[X] \wedge \text{Matrix}[X]\}$$

$$P_{\text{post}} : \{(D - \omega L)\underline{X} = (\omega U + (1 - \omega)D)\underline{X} + \omega \underline{b}e^T \\ \|\underline{X}e_r^T - \underline{X}e_{r-1}^T\| < \varepsilon\}$$

B.2.4. Richardson Iteration

$$P_{\text{pre}} : \{\text{Input}[A] \wedge \text{Matrix}[A] \wedge \text{NonSingular}[A] \wedge \\ \text{Input}[b] \wedge \text{Vector}[b] \wedge \\ \text{Input}[\alpha] \wedge \text{Scalar}[\alpha] \\ \text{FirstColumnInput}[X] \wedge \text{Matrix}[X] \wedge \\ \text{Output}[R] \wedge \text{Matrix}[R]\}$$

$$P_{\text{post}} : \{\alpha R = X(\underline{I} - \underline{J}) \\ R = A\underline{X} - \underline{b}e^T \\ \|\underline{R}e_r^T\| < \varepsilon\}$$

Bibliography

- [1] Richard Barrett, Michael W. Berry, Tony F. Chan, James Demmel, June Donato, Jack Dongarra, Victor Eijkhout, Roldan Pozo, Charles Romine, and Henk A. van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, volume 43. Siam, 1994.
- [2] Paolo Bientinesi. *Mechanical Derivation and Systematic Analysis of Correct Linear Algebra Algorithms*. Ph.D. Thesis, Department of Computer Sciences, University of Texas at Austin, July 2006.
- [3] Jack Edmonds. Maximum Matching and a Polyhedron with 0, 1-Vertices. *J. Res. Nat. Bur. Standards B*, 69(1965):125–130, 1965.
- [4] Victor Eijkhout, Paolo Bientinesi, and Robert van de Geijn. FLAME Derivation of CG Variants. Technical report, Texas Advanced Computing Center, The University of Texas at Austin, 2010.
- [5] Victor Eijkhout, Paolo Bientinesi, and Robert van de Geijn. Proof-Driven Derivation of Krylov Solver Libraries. Technical report, Aachen Institute for Computational Engineering Science, RWTH Aachen, June 2010. Technical Report AICES-2010/06-3.
- [6] Diego Fabregat-Traver. *Knowledge-Based Automatic Generation of Linear Algebra Algorithms and Code*. Ph.D. Thesis, RWTH Aachen, April 2014.
- [7] Diego Fabregat-Traver and Paolo Bientinesi. Knowledge-Based Automatic Generation of Partitioned Matrix Expressions. In Vladimir Gerdts, Wolfram Koepf, Ernst Mayr, and Evgenii Vorozhtsov, editors, *Computer Algebra in Scientific Computing*, volume 6885 of *Lecture Notes in Computer Science*, pages 144–157, Heidelberg, 2011. Springer.
- [8] David Gries and Fred B. Schneider. *A Logical Approach to Discrete Math*. Texts and Monographs in Computer Science. Springer Verlag, 1992.
- [9] Magnus Rudolph Hestenes and Eduard Stiefel. *Methods of Conjugate Gradients for Solving Linear Systems*. 1952.
- [10] Steven S. Muchnick. *Advanced Compiler Design and Implementation*. Morgan Kaufmann, 1997.

Bibliography

- [11] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. 2000.
- [12] G. W. Stewart. *Matrix Algorithms I: Basic Decompositions*. SIAM, Philadelphia, 1998.
- [13] Henk A. van der Vorst. *Iterative Krylov Methods for Large Linear Systems*, volume 13. Cambridge University Press, 2003.